

Newton-Stein Method: An optimization method for GLMs via Stein's Lemma

Murat A. Erdogdu*

Abstract

We consider the problem of efficiently computing the maximum likelihood estimator in *Generalized Linear Models* (GLMs) when the number of observations is much larger than the number of coefficients ($n \gg p \gg 1$). In this regime, optimization algorithms can immensely benefit from approximate second order information. We propose an alternative way of constructing the curvature information by formulating it as an estimation problem and applying a *Stein-type lemma*, which allows further improvements through sub-sampling and eigenvalue thresholding. Our algorithm enjoys fast convergence rates, resembling that of second order methods, with modest per-iteration cost. We provide its convergence analysis for the general case where the rows of the design matrix are samples from a sub-gaussian distribution. We show that the convergence has two phases, a quadratic phase followed by a linear phase. Finally, we empirically demonstrate that our algorithm achieves the highest performance compared to various algorithms on several datasets.

1 Introduction

Generalized Linear Models (GLMs) play a crucial role in numerous statistical and machine learning problems. GLMs formulate the natural parameter in exponential families as a linear model and provide a miscellaneous framework for statistical methodology and supervised learning tasks. Celebrated examples include linear, logistic, multinomial regressions and applications to graphical models [NB72, MN89, KF09].

In this paper, we focus on how to solve the maximum likelihood problem efficiently in the GLM setting when the number of observations n is much larger than the dimension of the coefficient vector p , i.e., $n \gg p \gg 1$. GLM optimization task is typically expressed as a minimization problem where the objective function is the negative log-likelihood that is denoted by $\ell(\beta)$ where $\beta \in \mathbb{R}^p$ is the coefficient vector. Many optimization algorithms are available for such minimization problems [Bis95, BV04, Nes04]. However, only a few uses the special structure of GLMs. In this paper, we consider updates that are specifically designed for GLMs, which are of the form

$$\beta \leftarrow \beta - \gamma \mathbf{Q} \nabla_{\beta} \ell(\beta), \quad (1.1)$$

where γ is the step size and \mathbf{Q} is a scaling matrix which provides curvature information.

For the updates of the form Eq. (1.1), the performance of the algorithm is mainly determined by the scaling matrix \mathbf{Q} . Classical *Newton's Method* (NM) and *Natural Gradient* (NG) descent can be

*Department of Statistics, Stanford University

recovered by simply taking \mathbf{Q} to be the inverse Hessian and the inverse Fisher’s information at the current iterate, respectively [Ama98, Nes04]. Second order methods may achieve quadratic convergence rate, yet they suffer from excessive cost of computing the scaling matrix at every iteration. On the other hand, if we take \mathbf{Q} to be the identity matrix, we recover the simple *Gradient Descent* (GD) method which has a linear convergence rate. Although GD’s convergence rate is slow compared to that of second order methods such as NM, modest per-iteration cost makes it practical for large-scale optimization.

The trade-off between the convergence rate and per-iteration cost has been extensively studied [Bis95, BV04, Nes04]. In $n \gg p$ regime, the main objective is to construct a scaling matrix \mathbf{Q} that is computational feasible and provides sufficient curvature information. For this purpose, several Quasi-Newton methods have been proposed [Bis95, Nes04]. Updates given by Quasi-Newton methods satisfy an equation which is often called *Quasi-Newton relation*. A well-known member of this class of algorithms is the *Broyden-Fletcher-Goldfarb-Shanno* (BFGS) algorithm [Bro70, Fle70, Gol70, Sha70].

In this paper, we propose an algorithm which utilizes the special structure of GLMs by relying on a Stein-type lemma [Ste81]. It attains fast convergence rates with low per-iteration cost. We call our algorithm *Newton-Stein Method* which we abbreviate as *NewSt*. Our contributions can be summarized as follows:

- We recast the problem of constructing a scaling matrix as an estimation problem and apply a Stein-type lemma along with sub-sampling to form a computationally feasible \mathbf{Q} .
- Our formulation allows further improvements through sub-sampling techniques and eigenvalue thresholding.
- Newton method’s $\mathcal{O}(np^2 + p^3)$ per-iteration cost is replaced by $\mathcal{O}(np + p^2)$ per-iteration cost and a one-time $\mathcal{O}(n|S|^2)$ cost, where $|S|$ is the sub-sample size.
- Assuming that the rows of the design matrix are i.i.d. and have bounded support (or sub-gaussian), and denoting the iterates of Newton-Stein Method by $\{\hat{\beta}^t\}_t$, we prove a bound of the form

$$\|\hat{\beta}^{t+1} - \beta_*\|_2 \leq \tau_1 \|\hat{\beta}^t - \beta_*\|_2 + \tau_2 \|\hat{\beta}^t - \beta_*\|_2^2, \quad (1.2)$$

where β_* is the minimizer and τ_1, τ_2 are the convergence coefficients. The above bound implies that the convergence starts with a quadratic phase and transitions into linear as the iterate gets closer to β_* .

- We demonstrate the performance of NewSt on four datasets by comparing it to commonly used algorithms.

The rest of the paper is organized as follows: Section 1.1 surveys the related work and Section 1.2 introduces the notations used throughout the paper. Section 2 briefly discusses the GLM framework and its relevant properties. In Section 3, we introduce Newton-Stein method, develop its intuition, and discuss the computational aspects. Section 4 covers the theoretical results and in Section 4.4 we discuss how to choose the algorithm parameters. Section 5 provides the empirical results where we compare the proposed algorithm with several other methods on four datasets. Finally, in Section 6, we conclude with a brief discussion along with a few open questions.

1.1 Related work

There are numerous optimization techniques that can be used to find the maximum likelihood estimator in GLMs. For moderate values of n and p , classical second order methods such as NM are commonly used. In large-scale problems, data dimensionality is the main factor while choosing the right optimization method. Large-scale optimization has been extensively studied through online and batch methods. Online methods use a gradient (or sub-gradient) of a single, randomly selected observation to update the current iterate [RM51]. Their per-iteration cost is independent of n , but the convergence rate might be extremely slow. There are several extensions of the classical stochastic descent algorithms (SGD), providing significant improvement and/or stability [Bot10, DHS11, SRB13].

On the other hand, batch algorithms enjoy faster convergence rates, though their per-iteration cost may be prohibitive. In particular, second order methods enjoy quadratic convergence, but constructing the Hessian matrix generally requires excessive amount of computation. Many algorithms aim at forming an approximate, cost-efficient scaling matrix. In particular, this idea lies at the core of Quasi-Newton methods [Bis95, Nes04].

Another approach to construct an approximate Hessian makes use of sub-sampling techniques [Mar10, BCNN11, VP12, EM15]. Many contemporary learning methods rely on sub-sampling as it is simple and it provides significant boost over the first order methods. Further improvements through conjugate gradient methods and Krylov sub-spaces are available. Sub-sampling can also be used to obtain an approximate solution, with certain large deviation guarantees [DLFU13].

There are many composite variants of the aforementioned methods, that mostly combine two or more techniques. Well-known composite algorithms are the combinations of sub-sampling and Quasi-Newton [SYG07, BHNS14], SGD and GD [FS12], NG and NM [LRF10], NG and low-rank approximation [LRMB08], sub-sampling and eigenvalue thresholding [EM15].

Lastly, algorithms that specialize on certain types of GLMs include coordinate descent methods for the penalized GLMs [FHT10] and trust region Newton methods [LWK08].

1.2 Notation

Let $[n] = \{1, 2, \dots, n\}$ and denote by $|S|$, the size of a set S . The gradient and the Hessian of f with respect to β are denoted by $\nabla_{\beta} f$ and $\nabla_{\beta}^2 f$, respectively. The j -th derivative of a function $f(w)$ is denoted by $f^{(j)}(w)$. For a vector x and a symmetric matrix \mathbf{X} , $\|x\|_2$ and $\|\mathbf{X}\|_2$ denote the ℓ_2 and spectral norms of x and \mathbf{X} , respectively. $\|x\|_{\psi_2}$ denotes the sub-gaussian norm, which will be defined later. S^{p-1} denotes the p -dimensional sphere. $\mathcal{P}_{\mathcal{C}}$ denotes the Euclidean projections onto the set \mathcal{C} , and $B_p(R) \subset \mathbb{R}^p$ denotes the ball of radius R . For a random variable x and density f , $x \sim f$ means that the distribution of x follows the density f . Multivariate Gaussian density with mean $\mu \in \mathbb{R}^p$ and covariance $\Sigma \in \mathbb{R}^{p \times p}$ is denoted as $\mathbf{N}_p(\mu, \Sigma)$. For random variables x, y , $d(x, y)$ and $\mathfrak{D}(x, y)$ denote probability metrics (to be explicitly defined later) measuring the distance between the distributions of x and y . $\mathcal{N}_{[]}(\cdots)$ and T_{ϵ} denote the bracketing number and ϵ -net.

2 Generalized Linear Models

Distribution of a random variable $y \in \mathbb{R}$ belongs to an exponential family with natural parameter $\eta \in \mathbb{R}$ if its density is

$$f(y|\eta) = e^{\eta y - \phi(\eta)} h(y),$$

where ϕ is the *cumulant generating function* and h is the *carrier density*. Let y_1, y_2, \dots, y_n be independent observations such that $\forall i \in [n], y_i \sim f(y_i|\eta_i)$. Denoting $\eta = (\eta_1, \dots, \eta_n)$, the joint likelihood can be written as

$$f(y_1, y_2, \dots, y_n|\eta) = \exp \left\{ \sum_{i=1}^n [y_i \eta_i - \phi(\eta_i)] \right\} \prod_{i=1}^n h(y_i). \quad (2.1)$$

We consider the problem of learning the maximum likelihood estimator in the above exponential family framework, where the vector $\eta \in \mathbb{R}^n$ is modeled through the linear relation,

$$\eta = \mathbf{X}\beta,$$

for some design matrix $\mathbf{X} \in \mathbb{R}^{n \times p}$ with rows $x_i \in \mathbb{R}^p$, and a coefficient vector $\beta \in \mathbb{R}^p$. This formulation is known as *Generalized Linear Models* (GLMs). The cumulant generating function ϕ determines the class of GLMs, i.e., for the ordinary least squares (OLS) $\phi(z) = z^2$ and for the logistic regression (LR) $\phi(z) = \log(1 + e^z)$.

Finding the maximum likelihood estimator in the above formulation is equivalent to minimizing the negative log-likelihood function $\ell(\beta)$,

$$\ell(\beta) = \frac{1}{n} \sum_{i=1}^n [\phi(\langle x_i, \beta \rangle) - y_i \langle x_i, \beta \rangle], \quad (2.2)$$

where $\langle x, \beta \rangle$ is the inner product between the vectors x and β . The relation to OLS and LR can be seen much easier by plugging in the corresponding $\phi(z)$ in Eq. (2.2). The gradient and the Hessian of $\ell(\beta)$ can be written as:

$$\nabla_{\beta} \ell(\beta) = \frac{1}{n} \sum_{i=1}^n [\phi^{(1)}(\langle x_i, \beta \rangle) x_i - y_i x_i], \quad \nabla_{\beta}^2 \ell(\beta) = \frac{1}{n} \sum_{i=1}^n \phi^{(2)}(\langle x_i, \beta \rangle) x_i x_i^T. \quad (2.3)$$

For a sequence of scaling matrices $\{\mathbf{Q}^t\}_{t>0} \in \mathbb{R}^{p \times p}$, we consider iterations of the form

$$\hat{\beta}^{t+1} \leftarrow \hat{\beta}^t - \gamma_t \mathbf{Q}^t \nabla_{\beta} \ell(\hat{\beta}^t)$$

where γ_t is the step size. The above iteration is our main focus, but with a new approach on how to compute the sequence of matrices $\{\mathbf{Q}^t\}_{t>0}$. We will formulate the problem of finding a scalable \mathbf{Q}^t as an estimation problem and apply a *Stein*-type lemma that provides us with a computationally efficient update.

3 Newton-Stein Method

Classical Newton-Raphson update is generally used for training GLMs. However, its per-iteration cost makes it impractical for large-scale optimization. The main bottleneck is the computation of the Hessian matrix that requires $\mathcal{O}(np^2)$ flops which is prohibitive when $n \gg p$. Numerous methods have been proposed to achieve NM's fast convergence rate while keeping the per-iteration cost manageable. To this end, a popular approach is to construct a scaling matrix \mathbf{Q}^t , which approximates the true Hessian at every iteration t .

Algorithm 1 Newton-Stein method

Input: $\hat{\beta}^0, r, \epsilon, \gamma$.

1. Set $t = 0$ and sub-sample a set of indices $S \subset [n]$ uniformly at random.
2. **Compute:** $\hat{\sigma}^2 = \lambda_{r+1}(\hat{\Sigma}_S)$, and $\zeta_r(\hat{\Sigma}_S) = \hat{\sigma}^2 \mathbf{I} + \operatorname{argmin}_{\operatorname{rank}(M)=r} \|\hat{\Sigma}_S - \hat{\sigma}^2 \mathbf{I} - M\|_F$.
3. **while** $\|\hat{\beta}^{t+1} - \hat{\beta}^t\|_2 \leq \epsilon$ **do**

$$\hat{\mu}_2(\hat{\beta}^t) = \frac{1}{n} \sum_{i=1}^n \phi^{(2)}(\langle x_i, \hat{\beta}^t \rangle), \quad \hat{\mu}_4(\hat{\beta}^t) = \frac{1}{n} \sum_{i=1}^n \phi^{(4)}(\langle x_i, \hat{\beta}^t \rangle),$$

$$\mathbf{Q}^t = \frac{1}{\hat{\mu}_2(\hat{\beta}^t)} \left[\zeta_r(\hat{\Sigma}_S)^{-1} - \frac{\hat{\beta}^t [\hat{\beta}^t]^T}{\hat{\mu}_2(\hat{\beta}^t)/\hat{\mu}_4(\hat{\beta}^t) + \langle \zeta_r(\hat{\Sigma}_S) \hat{\beta}^t, \hat{\beta}^t \rangle} \right],$$

$$\hat{\beta}^{t+1} = \mathcal{P}_{B_p(R)} \left(\hat{\beta}^t - \gamma \mathbf{Q}^t \nabla_{\beta} \ell(\hat{\beta}^t) \right),$$

$$t \leftarrow t + 1.$$
4. **end while**

Output: $\hat{\beta}^t$.

The task of constructing an approximate Hessian can be viewed as an estimation problem. Assuming that the rows of \mathbf{X} are i.i.d. random vectors, the Hessian of GLMs with cumulant generating function ϕ has the following form

$$[\mathbf{Q}^t]^{-1} = \frac{1}{n} \sum_{i=1}^n x_i x_i^T \phi^{(2)}(\langle x_i, \beta \rangle) \approx \mathbb{E}[x x^T \phi^{(2)}(\langle x, \beta \rangle)].$$

We observe that $[\mathbf{Q}^t]^{-1}$ is just a sum of i.i.d. matrices. Hence, the true Hessian is nothing but a sample mean estimator to its expectation. Another natural estimator would be the sub-sampled Hessian method suggested by [Mar10, BCNN11]. Therefore, our goal is to propose an estimator that is computationally efficient and well-approximates the true Hessian.

We use the following Stein-type proposition to find a more efficient estimator to the expectation of the Hessian.

Lemma 3.1 (Stein-type lemma). *Assume that $x \sim \mathcal{N}_p(0, \Sigma)$ and $\beta \in \mathbb{R}^p$ is a constant vector. Then for any function $f : \mathbb{R} \rightarrow \mathbb{R}$ that is twice “weakly” differentiable, we have*

$$\mathbb{E}[x x^T f(\langle x, \beta \rangle)] = \mathbb{E}[f(\langle x, \beta \rangle)] \Sigma + \mathbb{E}\left[f^{(2)}(\langle x, \beta \rangle)\right] \Sigma \beta \beta^T \Sigma. \quad (3.1)$$

The proof of Lemma 3.1 is given in Appendix. The right hand side of Eq.(3.1) is a rank-1 update to the first term. Hence, its inverse can be computed with $\mathcal{O}(p^2)$ cost. Quantities that change at each iteration are the ones that depend on β , i.e.,

$$\mu_2(\beta) = \mathbb{E}[\phi^{(2)}(\langle x, \beta \rangle)] \quad \text{and} \quad \mu_4(\beta) = \mathbb{E}[\phi^{(4)}(\langle x, \beta \rangle)].$$




Figure 1: The left plot demonstrates the accuracy of proposed Hessian estimation over different distributions. Number of observations is set to be $n = \mathcal{O}(p \log(p))$. The right plot shows the phase transition in the convergence rate of Newton-Stein method (NewSt). Convergence starts with a quadratic rate and transitions into linear. Plots are obtained using *Covtype* dataset.

Note that $\mu_2(\beta)$ and $\mu_4(\beta)$ are scalar quantities and can be estimated by their corresponding sample means $\hat{\mu}_2(\beta)$ and $\hat{\mu}_4(\beta)$ (explicitly defined at Step 3 of Algorithm 1) respectively, with only $\mathcal{O}(np)$ computation.

To complete the estimation task suggested by Eq. (3.1), we need an estimator for the covariance matrix Σ . A natural estimator is the sample mean where, we only use a sub-sample $S \subset [n]$ so that the cost is reduced to $\mathcal{O}(|S|p^2)$ from $\mathcal{O}(np^2)$. Sub-sampling based sample mean estimator is denoted by $\hat{\Sigma}_S = \sum_{i \in S} x_i x_i^T$, which is widely used in large-scale problems [Ver10]. We highlight the fact that Lemma 3.1 replaces NM's $\mathcal{O}(np^2)$ per-iteration cost with a one-time cost of $\mathcal{O}(np^2)$. We further use sub-sampling to reduce this one-time cost to $\mathcal{O}(|S|p^2)$.

In general, important curvature information is contained in the largest few spectral features [EM15]. For a given threshold r , we take the largest r eigenvalues of the sub-sampled covariance estimator, setting rest of them to $(r+1)$ -th eigenvalue. This operation helps denoising and would only take $\mathcal{O}(rp^2)$ computation. Step 2 of Algorithm 1 performs this procedure.

Inverting the constructed Hessian estimator can make use of the low-rank structure several times. First, notice that the updates in Eq. (3.1) are based on rank-1 matrix additions. Hence, we can simply use a matrix inversion formula to derive an explicit equation (See \mathbf{Q}^t in Step 3 of Algorithm 1). This formulation would impose another inverse operation on the covariance estimator. Since the covariance estimator is also based on rank- r approximation, one can utilize the low-rank inversion formula again. We emphasize that this operation is performed once. Therefore, instead of NM's per-iteration cost of $\mathcal{O}(p^3)$ due to inversion, Newton-Stein method (NewSt) requires $\mathcal{O}(p^2)$ per-iteration and a one-time cost of $\mathcal{O}(rp^2)$. Assuming that NewSt and NM converge in T_1 and T_2 iterations respectively, the overall complexity of NewSt is $\mathcal{O}(npT_1 + p^2T_1 + (|S| + r)p^2) \approx \mathcal{O}(npT_1 + p^2T_1 + |S|p^2)$ whereas that of NM is $\mathcal{O}(np^2T_2 + p^3T_2)$.

Even though Proposition 3.1 assumes that the covariates are multivariate Gaussian random

vectors, in Section 4, the only assumption we make on the covariates is either bounded support or sub-gaussianity, both of which cover a wide class of random variables including Bernoulli, elliptical distributions, bounded variables etc. The left plot of Figure 1 shows that the estimation is accurate for many distributions. This is a consequence of the fact that the proposed estimator in Eq. (3.1) relies on the distribution of x only through inner products of the form $\langle x, v \rangle$, which in turn results in approximate normal distribution due to the central limit theorem. We will discuss this in details in Section 4.

The convergence rate of NewSt has two phases. Convergence starts quadratically and transitions into linear rate when it gets close to the true minimizer. The phase transition behavior can be observed through the right plot in Figure 1. This is a consequence of the bound provided in Eq. 1.2, which is the main result of our theorems stated in Section 4.

4 Theoretical results

We start by introducing the terms that will appear in the theorems. Then we will provide two technical results on bounded and sub-gaussian covariates. The proofs of the theorems are technical and provided in Appendix.

4.1 Preliminaries

Hessian estimation described in the previous section relies on a Gaussian approximation. For theoretical purposes, we use the following probability metric to quantify the gap between the distribution of x_i 's and that of a normal vector.

Definition 1. Given a family of functions \mathcal{H} , and random vectors $x, y \in \mathbb{R}^p$, for \mathcal{H} and any $h \in \mathcal{H}$, define

$$d_{\mathcal{H}}(x, y) = \sup_{h \in \mathcal{H}} d_h(x, y) \quad \text{where} \quad d_h(x, y) = |\mathbb{E}[h(x)] - \mathbb{E}[h(y)]|.$$

Many probability metrics can be expressed as above by choosing a suitable function class \mathcal{H} . Examples include *Total Variation* (TV), *Kolmogorov* and *Wasserstein* metrics [GS02, CGS10]. Based on the second and the fourth derivatives of the cumulant generating function, we define the following function classes:

$$\begin{aligned} \mathcal{H}_1 &= \left\{ h(x) = \phi^{(2)}(\langle x, \beta \rangle) : \beta \in B_p(R) \right\}, & \mathcal{H}_2 &= \left\{ h(x) = \phi^{(4)}(\langle x, \beta \rangle) : \beta \in B_p(R) \right\}, \\ \mathcal{H}_3 &= \left\{ h(x) = \langle v, x \rangle^2 \phi^{(2)}(\langle x, \beta \rangle) : \beta \in B_p(R), \|v\|_2 = 1 \right\}, \end{aligned}$$

where $B_p(R) \in \mathbb{R}^p$ is the ball of radius R . Exact calculation of such probability metrics are often difficult. The general approach is to upper bound the distance by a more intuitive metric. In our case, we observe that $d_{\mathcal{H}_j}(x, y)$ for $j = 1, 2, 3$, can be easily upper bounded by $d_{\text{TV}}(x, y)$ up to a scaling constant, when the covariates have bounded support.

We will further assume that the covariance matrix follows r -spiked model, i.e.,

$$\Sigma = \sigma^2 I + \sum_{i=1}^r \theta_i u_i u_i^T,$$

which is commonly encountered in practice [BS06]. The first r eigenvalues of the covariance matrix are large and the rest are small and equal to each other. Small eigenvalues of Σ (denoted by σ^2), can be thought of as the noise component.

4.2 Bounded covariates

We have the following per-step bound for the iterates generated by NewSt, when the covariates are supported on a ball.

Theorem 4.1. *Assume that the covariates x_1, x_2, \dots, x_n are i.i.d. random vectors supported on a ball of radius \sqrt{K} with*

$$\mathbb{E}[x_i] = 0 \quad \text{and} \quad \mathbb{E}[x_i x_i^T] = \Sigma,$$

where Σ follows the r -spiked model. Further assume that the cumulant generating function ϕ has bounded 2nd-5th derivatives and that R is the radius of the projection $\mathcal{P}_{B_p(R)}$. For $\{\hat{\beta}^t\}_{t \geq 0}$ given by the Newton-Stein method for $\gamma = 1$, define the event

$$\mathcal{E} = \left\{ \left| \mu_2(\hat{\beta}^t) + \mu_4(\hat{\beta}^t) \langle \Sigma \hat{\beta}^t, \hat{\beta}^t \rangle \right| > \xi, \quad \beta_* \in B_p(R) \right\} \quad (4.1)$$

for some positive constant ξ , and the optimal value β_* . If $n, |S|$ and p are sufficiently large, then there exist constants c, c_1, c_2 and κ depending on the radii K, R , $\mathbb{P}(\mathcal{E})$ and the bounds on $|\phi^{(2)}|$ and $|\phi^{(4)}|$ such that conditioned on the event \mathcal{E} , with probability at least $1 - c/p^2$, we have

$$\|\hat{\beta}^{t+1} - \beta_*\|_2 \leq \tau_1 \|\hat{\beta}^t - \beta_*\|_2 + \tau_2 \|\hat{\beta}^t - \beta_*\|_2^2, \quad (4.2)$$

where the coefficients τ_1 and τ_2 are deterministic constants defined as

$$\tau_1 = \kappa \mathfrak{D}(x, z) + c_1 \kappa \sqrt{\frac{p}{\min\{p/\log(p)|S|, n/\log(n)\}}}, \quad \tau_2 = c_2 \kappa, \quad (4.3)$$

and $\mathfrak{D}(x, z)$ is defined as

$$\mathfrak{D}(x, z) = \|\Sigma\|_2 d_{\mathcal{H}_1}(x, z) + \|\Sigma\|_2^2 R^2 d_{\mathcal{H}_2}(x, z) + d_{\mathcal{H}_3}(x, z), \quad (4.4)$$

for a multivariate Gaussian random variable z with the same mean and covariance as x_i 's.

The bound in Eq. (4.2) holds with high probability, and the coefficients τ_1 and τ_2 are deterministic constants which will describe the convergence behavior of the Newton-Stein method. Observe that the coefficient τ_1 is sum of two terms: $\mathfrak{D}(x, z)$ measures how accurate the Hessian estimation is, and the second term depends on the sub-sample size and the data dimensions.

Theorem 4.1 shows that the convergence of Newton-Stein method can be upper bounded by a compositely converging sequence, that is, the squared term will dominate at first giving a quadratic rate, then the convergence will transition into a linear phase as the iterate gets close to the optimal value. The coefficients τ_1 and τ_2 govern the linear and quadratic terms, respectively. The effect of sub-sampling appears in the coefficient of linear term. In theory, there is a threshold for the sub-sampling size $|S|$, namely $\mathcal{O}(n/\log(n))$, beyond which further sub-sampling has no effect. The transition point between the quadratic and the linear phases is determined by the sub-sampling size and the properties of the data. The phase transition behavior can be observed through the right plot in Figure 1.

Using the above theorem, we state the following corollary.

Corollary 4.2. Assume that the assumptions of Theorem 4.1 hold. For a constant $\delta \geq \mathbb{P}(\mathcal{E}^C)$, a tolerance ϵ satisfying

$$\epsilon \geq 20R \{c/p^2 + \delta\},$$

and for an iterate satisfying $\mathbb{E}[\|\hat{\beta}^t - \beta_*\|_2] > \epsilon$, the iterates of Newton-Stein method will satisfy,

$$\mathbb{E}[\|\hat{\beta}^{t+1} - \beta_*\|_2] \leq \tilde{\tau}_1 \mathbb{E}[\|\hat{\beta}^t - \beta_*\|_2] + \tau_2 \mathbb{E}[\|\hat{\beta}^t - \beta_*\|_2^2],$$

where $\tilde{\tau}_1 = \tau_1 + 0.1$ and τ_1, τ_2 are as in Theorem 4.1.

The bound stated in the above corollary is an analogue of composite convergence (given in Eq. (4.2)) in expectation. Note that our results make strong assumptions on the derivatives of the cumulant generating function ϕ . We emphasize that these assumptions are valid for linear and logistic regressions. An example that does not fit in our scheme is *Poisson regression* with $\phi(z) = e^z$. However, we observed empirically that the algorithm still provides significant improvement.

The following theorem states a sufficient condition for the convergence of composite sequence.

Theorem 4.3. Let $\{\hat{\beta}^t\}_{t \geq 0}$ be a compositely converging sequence with convergence coefficients τ_1 and τ_2 as in Eq. (4.2) to the true minimizer β_* . Let the starting point satisfy $\|\hat{\beta}^0 - \beta_*\|_2 = \vartheta < (1 - \tau_1)/\tau_2$ and define $\Xi = \left(\frac{\tau_1 \vartheta}{1 - \tau_2 \vartheta}, \vartheta\right)$. Then the sequence of ℓ_2 -distances converges to 0. Further, the number of iterations to reach a tolerance of ϵ can be upper bounded by $\inf_{\xi \in \Xi} \mathcal{J}(\xi)$, where

$$\mathcal{J}(\xi) = \log_2 \left(\frac{\log(\delta(\tau_1/\xi + \tau_2))}{\log(\tau_1/\xi + \tau_2)\vartheta} \right) + \frac{\log(\epsilon/\xi)}{\log(\tau_1 + \tau_2\xi)}. \quad (4.5)$$

Above theorem gives an upper bound on the number of iterations until reaching a tolerance of ϵ . The first and second terms on the right hand side of Eq. (4.5) stem from the quadratic and linear phases, respectively.

4.3 Sub-gaussian covariates

In this section, we carry our analysis to the more general case, where the covariates are sub-gaussian vectors.

Theorem 4.4. Assume that x_1, x_2, \dots, x_n are i.i.d. sub-gaussian random vectors with sub-gaussian norm K such that

$$\mathbb{E}[x_i] = 0, \quad \mathbb{E}[\|x_i\|_2] = \mu \quad \text{and} \quad \mathbb{E}[x_i x_i^T] = \Sigma,$$

where Σ follows the r -spiked model. Further assume that the cumulant generating function ϕ is uniformly bounded and has bounded 2nd-5th derivatives and that R is the radius of the projection. For $\{\hat{\beta}^t\}_{t \geq 0}$ given by Newton-Stein method and the event \mathcal{E} in Eq. (4.1), if we have $n, |S|$ and p sufficiently large and

$$n^{0.2}/\log(n) \gtrsim p,$$

then there exist constants c_1, c_2, c_3, c_4 and κ depending on the eigenvalues of Σ , the radius R , μ , $\mathbb{P}(\mathcal{E})$ and the bounds on $\phi^{(2)}$ and $|\phi^{(4)}|$ such that conditioned on the event \mathcal{E} , with probability at least $1 - c_1 e^{-c_2 p}$,

$$\|\hat{\beta}^{t+1} - \beta_*\|_2 \leq \kappa \left\{ \mathfrak{D}(x, z) + c_3 \sqrt{\frac{p}{\min\{|S|, n^{0.2}/\log(n)\}}} + c_4 p^{1.5} \|\hat{\beta}^t - \beta_*\|_2 \right\} \|\hat{\beta}^t - \beta_*\|_2$$

where $\mathfrak{D}(x, z)$ defined as in Eq 4.4, for a Gaussian random variable z with the same mean and covariance as x_i 's.

The above theorem is more restrictive than Theorem 4.1. We require n to be much larger than the dimension p . Also note that a factor of $p^{1.5}$ appears in the coefficient of the quadratic term. We also notice that the threshold for sub-sample size reduces to $n^{0.2}/\log(n)$.

We have the following analogue of Corrolary 4.2.

Corollary 4.5. *Assume that the assumptions of Theorem 4.4 hold. For a constant $\delta \geq \mathbb{P}(\mathcal{E}^C)$, a tolerance ϵ satisfying*

$$\epsilon \geq 20R\sqrt{c_1 e^{-c_2 p} + \delta},$$

and for an iterate satisfying $\mathbb{E}[\|\hat{\beta}^t - \beta_*\|_2] > \epsilon$, the iterates of Newton-Stein method will satisfy,

$$\mathbb{E}[\|\hat{\beta}^{t+1} - \beta_*\|_2] \leq \tilde{\tau}_1 \mathbb{E}[\|\hat{\beta}^t - \beta_*\|_2] + \tau_2 \mathbb{E}[\|\hat{\beta}^t - \beta_*\|_2^2],$$

where $\tilde{\tau}_1 = \tau_1 + 0.1$ and τ_1, τ_2 are as in Theorem 4.4.

4.4 Algorithm parameters

Newton-Stein method takes three input parameters and for those, we suggest near-optimal choices based on our theoretical results.

- **Sub-sample size:** NewSt uses a subset of indices to approximate the covariance matrix Σ . Corollary 5.50 of [Ver10] proves that a sample size of $\mathcal{O}(p)$ is sufficient for sub-gaussian covariates and that of $\mathcal{O}(p \log(p))$ is sufficient for arbitrary distributions supported in some ball to estimate a covariance matrix by its sample mean estimator. In the regime we consider, $n \gg p$, we suggest to use a sample size of $|S| = \mathcal{O}(p \log(p))$.
- **Rank:** Many methods have been suggested to improve the estimation of covariance matrix and almost all of them rely on the concept of *shrinkage* [CCS10, DGJ13]. Eigenvalue thresholding can be considered as a shrinkage operation which will retain only the important second order information. Choosing the rank threshold r can be simply done on the sample mean estimator of Σ . After obtaining the sub-sampled estimate of the mean, one can either plot the spectrum and choose manually or use an optimal technique from [DG13].
- **Step size:** Step size choices of NewSt are quite similar to Newton's method (i.e., See [BV04]). The main difference comes from the eigenvalue thresholding. If the data follows the r -spiked

model, the optimal step size will be close to 1 if there is no sub-sampling. However, due to fluctuations resulting from sub-sampling, we suggest the following step size choice for NewSt :

$$\gamma = \frac{2}{1 + \frac{\hat{\sigma}^2 - \mathcal{O}(\sqrt{p/|S|})}{\hat{\sigma}^2}}. \quad (4.6)$$

This formula yields a step size larger than 1. A detailed discussion can be found in Section E in Appendix.

5 Experiments

In this section, we validate the performance of NewSt through extensive numerical studies. We experimented on two commonly used GLM optimization problems, namely, *Logistic Regression* (LR) and *Linear Regression* (OLS). LR minimizes Eq. (2.2) for the logistic function $\phi(z) = \log(1 + e^z)$, whereas OLS minimizes the same equation for $\phi(z) = z^2$. In the following, we briefly describe the algorithms that are used in the experiments:

- **Newton’s Method** (NM) uses the inverse Hessian evaluated at the current iterate, and may achieve quadratic convergence. NM steps require $\mathcal{O}(np^2 + p^3)$ computation which makes it impractical for large-scale datasets.
- **Broyden-Fletcher-Goldfarb-Shanno** (BFGS) forms a curvature matrix by cultivating the information from the iterates and the gradients at each iteration. Under certain assumptions, the convergence rate is locally super-linear and the per-iteration cost is comparable to that of first order methods.
- **Limited Memory BFGS** (L-BFGS) is similar to BFGS, and uses only the recent few iterates to construct the curvature matrix, gaining significant performance in terms of memory usage.
- **Gradient Descent** (GD) update is proportional to the negative of the full gradient evaluated at the current iterate. Under smoothness assumptions, GD achieves a linear convergence rate, with $\mathcal{O}(np)$ per-iteration cost.
- **Accelerated Gradient Descent** (AGD) is proposed by Nesterov [Nes83], which improves over the gradient descent by using a momentum term. Performance of AGD strongly depends of the smoothness of the function.

For all the algorithms, we use a constant step size that provides the fastest convergence. Sub-sample size $|S|$, rank r and the constant step-size γ for NewSt is selected by following the guidelines described in Section 4.4. The rank threshold r (which is an input to the algorithm) is specified on the title each plot.

5.1 Simulations with synthetic sub-gaussian datasets

Synthetic datasets, S3 and S20 are generated through a multivariate Gaussian distribution where the covariance matrix follows r -spiked model, i.e., $r = 3$ for S3 and $r = 20$ for S20. To generate the covariance matrix, we first generate a random orthogonal matrix, say \mathbf{M} . Next, we generate a

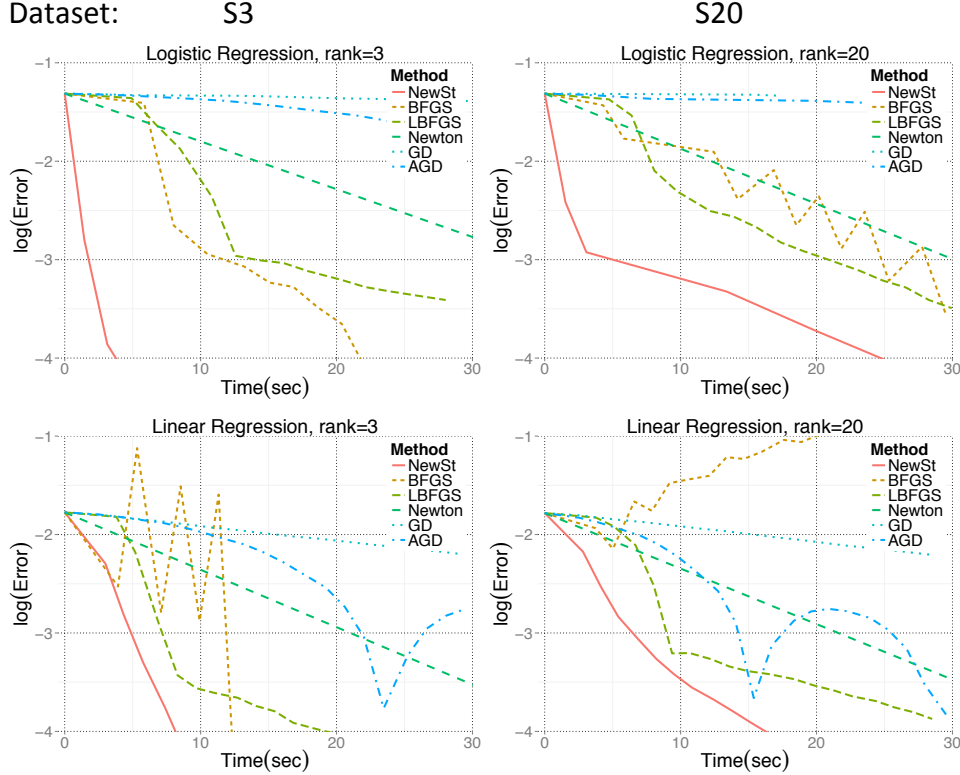


Figure 2: Performance of various optimization methods on synthetic datasets. Red straight line represents the proposed method NewSt . Algorithm parameters including the rank threshold is selected by the guidelines described in Section 4.4.

diagonal matrix $\mathbf{\Lambda}$ that contains the eigenvalues, i.e., the first r diagonal entries are chosen to be large, and rest of them are equal to 1. Then, we let $\mathbf{\Sigma} = \mathbf{M}\mathbf{\Lambda}\mathbf{M}^T$. For dimensions of the datasets, see Table 1. We also emphasize that the data dimensions are chosen so that Newton’s method still does well.

The simulation results are summarized in Figure 2. Further details regarding the experiments can be found in Table 2 in Appendix F. We observe that NewSt provides a significant improvement over the classical techniques.

Observe that the convergence rate of NewSt has a clear phase transition point in the top left plot in Figure 2. As argued earlier, this point depends on various factors including sub-sampling size $|S|$ and data dimensions n, p , the rank threshold r and structure of the covariance matrix. The prediction of the phase transition point is an interesting line of research. However, our convergence guarantees are conservative and cannot be used for this purpose.

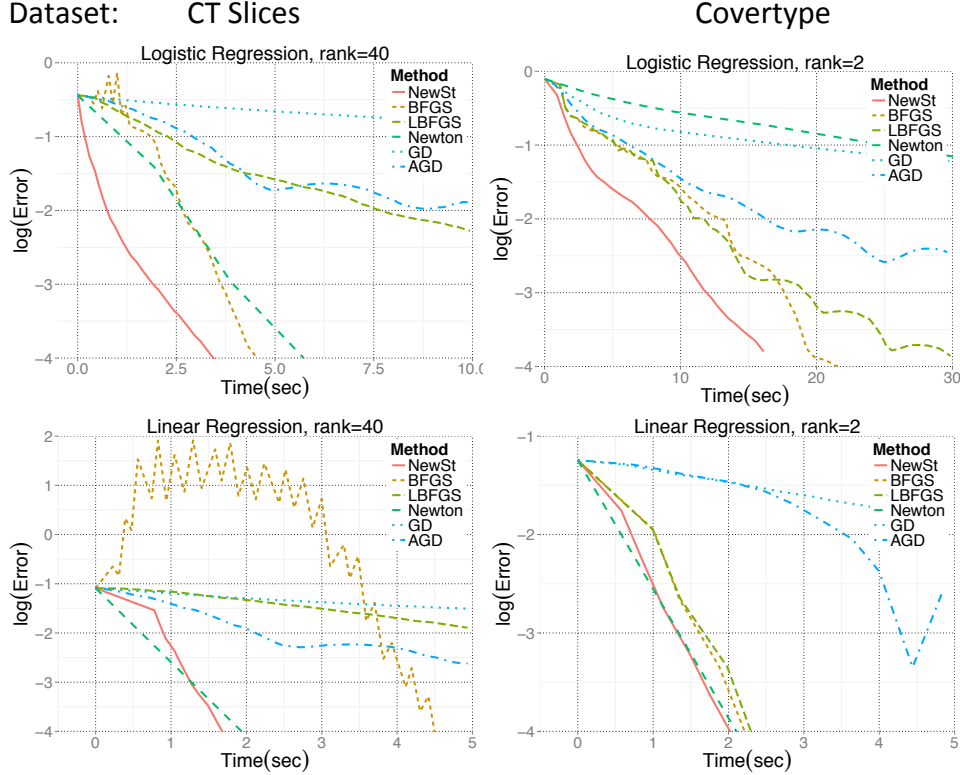


Figure 3: Performance of various optimization methods on real datasets. Red straight line represents the proposed method NewSt. Algorithm parameters including the rank threshold is selected by the guidelines described in Section 4.4.

5.2 Experiments with Real datasets

We experimented on two real datasets where the datasets are downloaded from UCI repository [Lic13]. Both datasets satisfy $n \gg p$, but we highlight the difference between the proportions of dimensions n/p . See Table 1 for details.

We observe that Newton-Stein method performs better than classical methods on real datasets as well. More specifically, the methods that come closer to NewSt is Newton’s method for moderate n and p and BFGS when n is large.

The optimal step-size for NewSt will typically be larger than 1 which is mainly due to eigenvalue thresholding operation. This feature is desirable if one is able to obtain a large step-size that provides convergence. In such cases, the convergence is likely to be faster, yet more unstable compared to the smaller step size choices. We observed that similar to other second order algorithms, NewSt is also susceptible to the step size selection. If the data is not well-conditioned, and the sub-sample size is not sufficiently large, algorithm might have poor performance. This is mainly because the sub-sampling operation is performed only once at the beginning. Therefore, it might be good in practice to sub-sample once in every few iterations.

Dataset	n	p	Reference, UCI repo [Lic13]
CT slices	53500	386	[GKS ⁺ 11]
Covertypes	581012	54	[BD99]
S3	500000	300	3-spiked model, [DGJ13]
S20	500000	300	20-spiked model, [DGJ13]

Table 1: Datasets used in the experiments.

6 Discussion

In this paper, we proposed an efficient algorithm for training GLMs. We called our algorithm Newton-Stein Method (NewSt) as it takes a Newton update at each iteration relying on a Stein-type lemma. The algorithm requires a one time $\mathcal{O}(|S|p^2)$ cost to estimate the covariance structure and $\mathcal{O}(np)$ per-iteration cost to form the update equations. We observe that the convergence of NewSt has a phase transition from quadratic rate to linear. This observation is justified theoretically along with several other guarantees for the sub-gaussian covariates such as per-step convergence bounds, conditions for convergence, etc. Parameter selection guidelines of NewSt are based on our theoretical results. Our experiments show that NewSt provides significant improvement over several optimization methods.

Relaxing some of the theoretical constraints is an interesting line of research. In particular, strong assumptions on the cumulant generating functions might be loosened. Another interesting direction is to determine when the phase transition point occurs, which would provide a better understanding of the effects of sub-sampling and rank threshold.

Acknowledgements

The author is grateful to Mohsen Bayati and Andrea Montanari for stimulating conversations on the topic of this work. The author would like to thank Bhaswar B. Bhattacharya and Qingyuan Zhao for carefully reading this article and providing valuable feedback.

References

- [Ama98] Shun-Ichi Amari, *Natural gradient works efficiently in learning*, Neural computation **10** (1998), no. 2, 251–276.
- [BCNN11] Richard H Byrd, Gillian M Chin, Will Neveitt, and Jorge Nocedal, *On the use of stochastic hessian information in optimization methods for machine learning*, SIAM Journal on Optimization **21** (2011), no. 3, 977–995.
- [BD99] Jock A Blackard and Denis J Dean, *Comparative accuracies of artificial neural networks and discriminant analysis in predicting forest cover types from cartographic variables*, Computers and electronics in agriculture (1999), 131–151.
- [BHNS14] Richard H Byrd, SL Hansen, Jorge Nocedal, and Yoram Singer, *A stochastic quasi-newton method for large-scale optimization*, arXiv preprint arXiv:1401.7020 (2014).

- [Bis95] Christopher M. Bishop, *Neural networks for pattern recognition*, Oxford University Press, Inc., NY, USA, 1995.
- [Bot10] L  on Bottou, *Large-scale machine learning with stochastic gradient descent*, COMPSTAT, 2010, pp. 177–186.
- [Bro70] Charles G Broyden, *The convergence of a class of double-rank minimization algorithms 2. the new algorithm*, IMA Journal of Applied Mathematics **6** (1970), no. 3, 222–231.
- [BS06] Jinho Baik and Jack W Silverstein, *Eigenvalues of large sample covariance matrices of spiked population models*, Journal of Multivariate Analysis **97** (2006), no. 6, 1382–1408.
- [BV04] Stephen Boyd and Lieven Vandenberghe, *Convex optimization*, Cambridge University Press, New York, NY, USA, 2004.
- [CCS10] Jian-Feng Cai, Emmanuel J Cand  s, and Zuowei Shen, *A singular value thresholding algorithm for matrix completion*, SIAM Journal on Optimization **20** (2010), no. 4, 1956–1982.
- [CGS10] Louis HY Chen, Larry Goldstein, and Qi-Man Shao, *Normal approximation by Steins method*, Springer Science, 2010.
- [DE15] Lee H Dicker and Murat A Erdogdu, *Flexible results for quadratic forms with applications to variance components estimation*, arXiv preprint arXiv:1509.04388 (2015).
- [DG13] David L Donoho and Matan Gavish, *The optimal hard threshold for singular values is $4/\sqrt{3}$* , arXiv:1305.5870 (2013).
- [DGJ13] David L Donoho, Matan Gavish, and Iain M Johnstone, *Optimal shrinkage of eigenvalues in the spiked covariance model*, arXiv preprint arXiv:1311.0851 (2013).
- [DHS11] John Duchi, Elad Hazan, and Yoram Singer, *Adaptive subgradient methods for online learning and stochastic optimization*, J. Mach. Learn. Res. **12** (2011), 2121–2159.
- [DLFU13] Paramveer Dhillon, Yichao Lu, Dean P Foster, and Lyle Ungar, *New subsampling algorithms for fast least squares regression*, Advances in Neural Information Processing Systems 26, 2013, pp. 360–368.
- [EM15] Murat A Erdogdu and Andrea Montanari, *Convergence rates of sub-sampled newton methods*, arXiv preprint arXiv:1508.02810 (2015).
- [FHT10] Jerome Friedman, Trevor Hastie, and Rob Tibshirani, *Regularization paths for generalized linear models via coordinate descent*, Journal of statistical software **33** (2010), no. 1, 1.
- [Fle70] Roger Fletcher, *A new approach to variable metric algorithms*, The computer journal **13** (1970), no. 3, 317–322.
- [FS12] Michael P Friedlander and Mark Schmidt, *Hybrid deterministic-stochastic methods for data fitting*, SIAM Journal on Scientific Computing **34** (2012), no. 3, A1380–A1405.

- [GKS⁺11] Franz Graf, Hans-Peter Kriegel, Matthias Schubert, Sebastian Pölsterl, and Alexander Cavallaro, *2d image registration in ct images using radial image descriptors*, MICCAI 2011, Springer, 2011, pp. 607–614.
- [Gol70] Donald Goldfarb, *A family of variable-metric methods derived by variational means*, Mathematics of computation **24** (1970), no. 109, 23–26.
- [GS02] Alison L Gibbs and Francis Edward Su, *On choosing and bounding probability metrics*, ISR **70** (2002), no. 3, 419–435.
- [KF09] Daphne Koller and Nir Friedman, *Probabilistic graphical models: principles and techniques*, MIT press, 2009.
- [Lic13] M. Lichman, *UCI machine learning repository*, 2013.
- [LRF10] Nicolas Le Roux and Andrew W Fitzgibbon, *A fast natural newton method*, ICML, 2010, pp. 623–630.
- [LRMB08] Nicolas Le Roux, Pierre-A Manzagol, and Yoshua Bengio, *Topmoumoute online natural gradient algorithm*, NIPS, 2008.
- [LWK08] Chih-J Lin, Ruby C Weng, and Sathya Keerthi, *Trust region newton method for logistic regression*, JMLR (2008).
- [Mar10] James Martens, *Deep learning via hessian-free optimization*, ICML, 2010, pp. 735–742.
- [MN89] Peter McCullagh and John A Nelder, *Generalized linear models*, vol. 2, Chapman and Hall London, 1989.
- [NB72] John A Nelder and R. Jacob Baker, *Generalized linear models*, Wiley Online Library, 1972.
- [Nes83] Yurii Nesterov, *A method for unconstrained convex minimization problem with the rate of convergence $O(1/k^2)$* , Doklady AN SSSR, vol. 269, 1983, pp. 543–547.
- [Nes04] ———, *Introductory lectures on convex optimization: A basic course*, vol. 87, Springer, 2004.
- [RM51] Herbert Robbins and Sutton Monro, *A stochastic approximation method*, Annals of mathematical statistics (1951).
- [Sha70] David F Shanno, *Conditioning of quasi-newton methods for function minimization*, Mathematics of computation **24** (1970), no. 111, 647–656.
- [SRB13] Mark Schmidt, Nicolas Le Roux, and Francis Bach, *Minimizing finite sums with the stochastic average gradient*, arXiv preprint arXiv:1309.2388 (2013).
- [Ste81] Charles M Stein, *Estimation of the mean of a multivariate normal distribution*, Annals of Statistics (1981), 1135–1151.
- [SYG07] Nicol Schraudolph, Jin Yu, and Simon Günter, *A stochastic quasi-newton method for online convex optimization*.

- [VdV00] Aad W Van der Vaart, *Asymptotic statistics*, vol. 3, Cambridge university press, 2000.
- [Ver10] Roman Vershynin, *Introduction to the non-asymptotic analysis of random matrices*, arXiv:1011.3027 (2010).
- [VP12] Oriol Vinyals and Daniel Povey, *Krylov Subspace Descent for Deep Learning*, AISTATS, 2012.

A Proof of Stein-type lemma

Proof of Lemma 3.1. The proof will follow from integration by parts over multivariate variables. Let $g(x)$ be the density of x , i.e.,

$$g(x) = (2\pi)^{-p/2} |\Sigma|^{-1/2} \exp \left\{ -\frac{1}{2} \langle \Sigma^{-1} x, x \rangle \right\},$$

and $xg(x)dx = -\Sigma dg(x)$. We write

$$\begin{aligned} \mathbb{E}[xx^T f(\langle x, \beta \rangle)] &= \int xx^T f(\langle x, \beta \rangle) g(x) dx \\ &= \Sigma \int -f(\langle x, \beta \rangle) dg(x) x^T, \\ &= \Sigma \left\{ \int f(\langle x, \beta \rangle) g(x) dx + \int \beta x^T \dot{f}(\langle x, \beta \rangle) g(x) dx \right\}, \\ &= \Sigma \left\{ \mathbb{E}[f(\langle x, \beta \rangle)] + \int \beta \beta^T \ddot{f}(\langle x, \beta \rangle) g(x) dx \Sigma \right\}, \\ &= \Sigma \left\{ \mathbb{E}[f(\langle x, \beta \rangle)] + \beta \beta^T \mathbb{E}[\ddot{f}(\langle x, \beta \rangle)] \Sigma \right\}, \\ &= \mathbb{E}[f(\langle x, \beta \rangle)] \Sigma + \mathbb{E}[\ddot{f}(\langle x, \beta \rangle)] \Sigma \beta \beta^T \Sigma, \end{aligned}$$

which is the desired result. \square

B Preliminary concentration inequalities

In this section, we provide concentration results that will be useful proving the main theorem. We start with some simple definitions on a special class of random variables.

Definition 2 (Sub-gaussian). *For a given constant K , a random variable $x \in \mathbb{R}$ is called sub-gaussian if it satisfies*

$$\mathbb{E}[|x|^m]^{1/m} \leq K \sqrt{m}, \quad m \geq 1.$$

Smallest such K is the sub-gaussian norm of x and it is denoted by $\|x\|_{\psi_2}$. Similarly, a random vector $y \in \mathbb{R}^p$ is a sub-gaussian vector if there exists a constant $K' > 0$ such that

$$\sup_{v \in S^{p-1}} \|\langle y, v \rangle\|_{\psi_2} \leq K'.$$

Definition 3 (Sub-exponential). *For a given constant K , a random variable $x \in \mathbb{R}$ is called sub-exponential if it satisfies*

$$\mathbb{E}[|x|^m]^{1/m} \leq Km, \quad m \geq 1,$$

Smallest such K is the sub-exponential norm of x and it is denoted by $\|x\|_{\psi_1}$. Similarly, a random vector $y \in \mathbb{R}^p$ is a sub-exponential vector if there exists a constant $K' > 0$ such that

$$\sup_{v \in S^{p-1}} \|\langle y, v \rangle\|_{\psi_1} \leq K'.$$

We state the following Lemmas from [Ver10] for the convenience of the reader (i.e., See Theorem 5.39 and the following remark for sub-gaussian distributions, and Theorem 5.44 for distributions with arbitrary support):

Lemma B.1 ([Ver10]). *Let S be an index set and $x_i \in \mathbb{R}^p$ for $i \in S$ be i.i.d. sub-gaussian random vectors with*

$$\mathbb{E}[x_i] = 0, \quad \mathbb{E}[x_i x_i^T] = \Sigma \quad \|x_i\|_{\psi_2} \leq K.$$

There exists absolute constants c, C depending only on the sub-gaussian norm K such that with probability $1 - 2e^{-ct^2}$,

$$\left\| \widehat{\Sigma}_S - \Sigma \right\|_2 \leq \max(\delta, \delta^2) \quad \text{where} \quad \delta = C \sqrt{\frac{p}{|S|}} + \frac{t}{\sqrt{|S|}}.$$

Remark 1. *We are interested in the case where $\delta < 1$, hence the right hand side becomes $\max(\delta, \delta^2) = \delta$. In most cases, we will simply let $t = \sqrt{p}$ and obtain a bound of order $\sqrt{p/|S|}$ on the right hand side. For this, we need $|S| = \mathcal{O}(C^2 p)$ which is a reasonable assumption in the regime we consider.*

The following lemma will be helpful to show a similar concentration result for the matrix $\zeta_r(\widehat{\Sigma}_S)$:

Lemma B.2. *Let the assumptions in Lemma B.1 hold. Further, assume that Σ follows r -spiked model. If $|S|$ is sufficiently large, then there exists absolute constants c, C depending only on the sub-gaussian norm K such that with probability $1 - 2e^{-cp}$,*

$$\left\| \zeta_r(\widehat{\Sigma}_S) - \widehat{\Sigma}_S \right\|_2 \leq C \sqrt{\frac{p}{|S|}}.$$

Proof. By the Weyl's inequality for the eigenvalues, we have

$$\left\| \zeta_r(\widehat{\Sigma}_S) - \widehat{\Sigma}_S \right\|_2 = \lambda_{r+1}(\widehat{\Sigma}_S) - \lambda_p(\widehat{\Sigma}_S) \leq 2 \left\| \widehat{\Sigma}_S - \Sigma \right\|_2.$$

Hence the result follows from the previous lemma for $t = \sqrt{p}$. □

Lemmas B.1 and B.2 are straightforward concentration results for the random matrices with i.i.d. sub-gaussian rows. We state their analogues for the the covariates sampled from arbitrary distributions with bounded support.

Lemma B.3 ([Ver10]). *Let S be an index set and $x_i \in \mathbb{R}^p$ for $i \in S$ be i.i.d. random vectors with*

$$\mathbb{E}[x_i] = 0, \quad \mathbb{E}[x_i x_i^T] = \Sigma, \quad \|x_i\|_2 \leq \sqrt{K} \text{ a.s.}$$

Then, for some absolute constant C , with probability $1 - pe^{-Ct^2}$, we have

$$\left\| \widehat{\Sigma}_S - \Sigma \right\|_2 \leq \max\left(\left\| \Sigma \right\|_2^{1/2} \delta, \delta^2\right) \quad \text{where} \quad \delta = t \sqrt{\frac{K}{|S|}}.$$

Remark 2. We will choose $t = \sqrt{3\log(p)/C}$ which will provide us with a probability of $1 - 1/p^2$. Therefore, if the sample size is sufficiently large, i.e.,

$$|S| \geq \frac{3K \log(p)}{C \|\Sigma\|_2} = \mathcal{O}(K \log(p) / \|\Sigma\|_2),$$

we can estimate the true covariance matrix quite well for arbitrary distributions with bounded support. In particular, with probability $1 - 1/p^2$, we obtain

$$\|\widehat{\Sigma}_S - \Sigma\|_2 \leq c \sqrt{\frac{\log(p)}{|S|}},$$

where $c = \sqrt{3K \|\Sigma\|_2 / C}$.

Lemma B.4. Let the assumptions in Lemma B.3 hold. Further, assume that Σ follows r -spiked model. If $|S|$ is sufficiently large, for $c = 2\sqrt{3K \|\Sigma\|_2 / C}$, with probability $1 - 1/p^2$, we have

$$\|\zeta_r(\widehat{\Sigma}_S) - \widehat{\Sigma}_S\|_2 \leq c \sqrt{\frac{\log(p)}{|S|}},$$

where C is an absolute constant.

Proof of Lemma B.4 is the same as that of Lemma B.2. Before proceeding, we note that the bounds given in Lemmas B.2 and B.4 also applies to $\|\zeta_r(\widehat{\Sigma}_S) - \Sigma\|_2$.

In the following, we will focus on the empirical processes and obtain more technical bounds for the approximate Hessian. To that extent, we provide some basic definitions that will be useful later in the proofs. For a more detailed discussion on the machinery used throughout next section, we refer interested reader to [VdV00].

Definition 4. On a metric space (X, d) , for $\epsilon > 0$, $T_\epsilon \subset X$ is called an ϵ -net over X if $\forall x \in X$, $\exists t \in T_\epsilon$ such that $d(x, t) \leq \epsilon$.

In the following, we will use L_1 distance between two functions f and g , namely $d(f, g) = \int |f - g|$. Note that the same distance definition can be carried to random variables as they are simply real measurable functions.

Definition 5. Given a function class \mathcal{F} , and any two functions l and u (not necessarily in \mathcal{F}), the bracket $[l, u]$ is the set of all $f \in \mathcal{F}$ such that $l \leq f \leq u$. A bracket satisfying $l \leq u$ and $\int |u - l| \leq \epsilon$ is called an ϵ -bracket in L_1 . The bracketing number $\mathcal{N}_{[]}(\epsilon, \mathcal{F}, L_1)$ is the minimum number of different ϵ -brackets needed to cover \mathcal{F} .

The preliminary tools presented in this section will be utilized to obtain the concentration results in Section C.

C Main lemmas

C.1 Concentration of covariates with bounded support

Lemma C.1. *Let $x_i \in \mathbb{R}^p$, for $i = 1, 2, \dots, n$, be i.i.d. random vectors supported on a ball of radius \sqrt{K} , with mean 0, and covariance matrix Σ . Also let $f : \mathbb{R} \rightarrow \mathbb{R}$ be a uniformly bounded function such that for some $B > 0$, we have $\|f\|_\infty < B$ and f is Lipschitz continuous with constant L . Then, there exist constants c_1, c_2, c_3 such that*

$$\mathbb{P} \left(\sup_{\beta \in B_n(R)} \left| \frac{1}{n} \sum_{i=1}^n f(\langle x_i, \beta \rangle) - \mathbb{E}[f(\langle x, \beta \rangle)] \right| > c_1 \sqrt{\frac{p \log(n)}{n}} \right) \leq c_2 e^{-c_3 p},$$

where the constants depend only on the bound B and radii R and \sqrt{K} .

Proof of Lemma C.1. We start by using the Lipschitz property of the function f , i.e., $\forall \beta, \beta' \in B_p(R)$,

$$\begin{aligned} \|f(\langle x, \beta \rangle) - f(\langle x, \beta' \rangle)\|_2 &\leq L \|x\|_2 \|\beta - \beta'\|_2, \\ &\leq L \sqrt{K} \|\beta - \beta'\|_2. \end{aligned}$$

Now let T_Δ be a Δ -net over $B_p(R)$. Then $\forall \beta \in B_p(R)$, $\exists \beta' \in T_\Delta$ such that right hand side of the above inequality is smaller than $\Delta L \sqrt{K}$. Then, we can write

$$\left| \frac{1}{n} \sum_{i=1}^n f(\langle x_i, \beta \rangle) - \mathbb{E}[f(\langle x, \beta \rangle)] \right| \leq \left| \frac{1}{n} \sum_{i=1}^n f(\langle x_i, \beta' \rangle) - \mathbb{E}[f(\langle x, \beta' \rangle)] \right| + 2\Delta L \sqrt{K}. \quad (\text{C.1})$$

By choosing

$$\Delta = \frac{\epsilon}{4L\sqrt{K}},$$

and taking supremum over the corresponding β sets on both sides, we obtain the following inequality

$$\sup_{\beta \in B_n(R)} \left| \frac{1}{n} \sum_{i=1}^n f(\langle x_i, \beta \rangle) - \mathbb{E}[f(\langle x, \beta \rangle)] \right| \leq \max_{\beta \in T_\Delta} \left| \frac{1}{n} \sum_{i=1}^n f(\langle x_i, \beta \rangle) - \mathbb{E}[f(\langle x, \beta \rangle)] \right| + \frac{\epsilon}{2}.$$

Now, since we have $\|f\|_\infty \leq B$ and for a fixed β and $i = 1, 2, \dots, n$, the random variables $f(\langle x_i, \beta \rangle)$ are i.i.d., by the Hoeffding's concentration inequality, we have

$$\mathbb{P} \left(\left| \frac{1}{n} \sum_{i=1}^n f(\langle x_i, \beta \rangle) - \mathbb{E}[f(\langle x, \beta \rangle)] \right| > \epsilon/2 \right) \leq 2 \exp \left(-\frac{n\epsilon^2}{2B^2} \right).$$

Combining Eq. (C.1), the above result together with the union bound, we easily obtain

$$\begin{aligned} \mathbb{P} \left(\sup_{\beta \in B_n(R)} \left| \frac{1}{n} \sum_{i=1}^n f(\langle x_i, \beta \rangle) - \mathbb{E}[f(\langle x, \beta \rangle)] \right| > \epsilon \right) &\leq \mathbb{P} \left(\max_{\beta \in T_\Delta} \left| \frac{1}{n} \sum_{i=1}^n f(\langle x_i, \beta \rangle) - \mathbb{E}[f(\langle x, \beta \rangle)] \right| > \epsilon/2 \right) \\ &\leq 2|T_\Delta| \exp \left(-\frac{n\epsilon^2}{2B^2} \right), \end{aligned}$$

where $\Delta = \epsilon/4L\sqrt{K}$.

Next, we apply Lemma G.5 and obtain that

$$|\mathcal{T}_\Delta| \leq \left(\frac{R\sqrt{p}}{\Delta} \right)^p = \left(\frac{R\sqrt{p}}{\epsilon/4L\sqrt{K}} \right)^p.$$

We require that the bound on the probability gets an exponential decay with rate $\mathcal{O}(p)$. Using Lemma G.6 with $a = 2B^2p/n$ and $b = 4LR\sqrt{Kp}$, we obtain that ϵ should be

$$\epsilon = \sqrt{\frac{B^2p}{n} \log \left(\frac{16L^2R^2Kn}{B^2} \right)} = \mathcal{O} \left(\sqrt{\frac{p \log(n)}{n}} \right),$$

which completes the proof. \square

In the following, we state the concentration results on functions of the form

$$x \rightarrow f(\langle x, \beta \rangle) \langle x, v \rangle^2.$$

Functions of this type form the summands of the Hessian matrix in GLMs.

Lemma C.2. *Let $x_i \in \mathbb{R}^p$, for $i = 1, \dots, n$, be i.i.d. random vectors supported on a ball of radius \sqrt{K} , with mean 0, covariance matrix Σ . Also let $f : \mathbb{R} \rightarrow \mathbb{R}$ be a uniformly bounded function such that for some $B > 0$, we have $\|f\|_\infty < B$ and f is Lipschitz continuous with constant L . Then, for $v \in S^{p-1}$, there exist constants c_1, c_2, c_3 such that*

$$\mathbb{P} \left(\sup_{\beta \in B_p(R)} \left| \frac{1}{n} \sum_{i=1}^n f(\langle x_i, \beta \rangle) \langle x_i, v \rangle^2 - \mathbb{E}[f(\langle x, \beta \rangle) \langle x, v \rangle^2] \right| > c_1 \sqrt{\frac{p \log(n)}{n}} \right) \leq c_2 e^{-c_3 p},$$

where the constants depend only on the bound B and radii R and \sqrt{K} .

Proof of Lemma C.2. As in the proof of Lemma C.1, we start by using the Lipschitz property of the function f , i.e., $\forall \beta, \beta' \in B_p(R)$,

$$\begin{aligned} \|f(\langle x, \beta \rangle) \langle x, v \rangle^2 - f(\langle x, \beta' \rangle) \langle x, v \rangle^2\|_2 &\leq L \|x\|_2^3 \|\beta - \beta'\|_2, \\ &\leq LK^{1.5} \|\beta - \beta'\|_2. \end{aligned}$$

For a net T_Δ , $\forall \beta \in B_p(R)$, $\exists \beta' \in T_\Delta$ such that right hand side of the above inequality is smaller than $\Delta L \sqrt{K}$. Then, we can write

$$\begin{aligned} \left| \frac{1}{n} \sum_{i=1}^n f(\langle x_i, \beta \rangle) \langle x_i, v \rangle^2 - \mathbb{E}[f(\langle x, \beta \rangle) \langle x, v \rangle^2] \right| &\leq \left| \frac{1}{n} \sum_{i=1}^n f(\langle x_i, \beta' \rangle) \langle x_i, v \rangle^2 - \mathbb{E}[f(\langle x, \beta' \rangle) \langle x, v \rangle^2] \right| \\ &\quad + 2\Delta L K^{1.5}. \end{aligned} \tag{C.2}$$

This time, we choose

$$\Delta = \frac{\epsilon}{4LK^{1.5}},$$

and take the supremum over the corresponding feasible β -sets on both sides,

$$\begin{aligned} & \sup_{\beta \in B_n(R)} \left| \frac{1}{n} \sum_{i=1}^n f(\langle x_i, \beta \rangle) \langle x_i, v \rangle^2 - \mathbb{E}[f(\langle x, \beta \rangle) \langle x, v \rangle^2] \right| \\ & \leq \max_{\beta \in \mathcal{T}_\Delta} \left| \frac{1}{n} \sum_{i=1}^n f(\langle x_i, \beta \rangle) \langle x_i, v \rangle^2 - \mathbb{E}[f(\langle x, \beta \rangle) \langle x, v \rangle^2] \right| + \frac{\epsilon}{2}. \end{aligned}$$

Now, since we have $\|f\|_\infty \leq B$ and for fixed β and v , $i = 1, 2, \dots, n$, $f(\langle x_i, \beta \rangle) \langle x_i, v \rangle^2$ are i.i.d. random variables. By the Hoeffding's concentration inequality, we have

$$\mathbb{P} \left(\left| \frac{1}{n} \sum_{i=1}^n f(\langle x_i, \beta \rangle) \langle x_i, v \rangle^2 - \mathbb{E}[f(\langle x, \beta \rangle) \langle x, v \rangle^2] \right| > \epsilon/2 \right) \leq 2 \exp \left(-\frac{n\epsilon^2}{2B^2K^2} \right).$$

Using Eq. (C.2) and the above result combined with the union bound, we easily obtain

$$\begin{aligned} & \mathbb{P} \left(\sup_{\beta \in B_n(R)} \left| \frac{1}{n} \sum_{i=1}^n f(\langle x_i, \beta \rangle) \langle x_i, v \rangle^2 - \mathbb{E}[f(\langle x, \beta \rangle) \langle x, v \rangle^2] \right| > \epsilon \right) \\ & \leq \mathbb{P} \left(\max_{\beta \in \mathcal{T}_\Delta} \left| \frac{1}{n} \sum_{i=1}^n f(\langle x_i, \beta \rangle) \langle x_i, v \rangle^2 - \mathbb{E}[f(\langle x, \beta \rangle) \langle x, v \rangle^2] \right| > \epsilon/2 \right) \\ & \leq 2|\mathcal{T}_\Delta| \exp \left(-\frac{n\epsilon^2}{2B^2K^2} \right), \end{aligned}$$

where $\Delta = \epsilon/4LK^{1.5}$. Using Lemma G.5, we have

$$|\mathcal{T}_\Delta| \leq \left(\frac{R\sqrt{p}}{\Delta} \right)^p = \left(\frac{R\sqrt{p}}{\epsilon/4LK^{1.5}} \right)^p.$$

As before, we require that the right hand side of above inequality gets a decay with rate $\mathcal{O}(p)$. Using Lemma G.6 with $a = 2B^2K^2p/n$ and $b = 4LRK^{1.5}\sqrt{p}$, we obtain that ϵ should be

$$\epsilon = \sqrt{\frac{B^2K^2p}{n} \log \left(\frac{16L^2R^2K^3n}{B^2} \right)} = \mathcal{O} \left(\sqrt{\frac{p \log(n)}{n}} \right),$$

which completes the proof. \square

C.2 Concentration of sub-gaussian covariates

In this section, we derive the analogues of the Lemmas C.1 and C.2 for sub-gaussian covariates. Note that the Lemmas in this section are more general in the sense that they also apply to the case where covariates have bounded support. However, the concentration coefficients are different (worse) compared to previous section.

Lemma C.3. *Let $x_i \in \mathbb{R}^p$, for $i = 1, \dots, n$, be i.i.d. sub-gaussian random vectors with mean 0, covariance matrix Σ and sub-gaussian norm K . Also let $f : \mathbb{R} \rightarrow \mathbb{R}$ be a uniformly bounded function*

such that for some $B > 0$, we have $\|f\|_\infty < B$ and f is Lipschitz continuous with constant L . Then, there exists absolute constants c_1, c_2, c_3 such that

$$\mathbb{P} \left(\sup_{\beta \in B_n(R)} \left| \frac{1}{n} \sum_{i=1}^n f(\langle x_i, \beta \rangle) - \mathbb{E}[f(\langle x, \beta \rangle)] \right| > c_1 \sqrt{\frac{p \log(n)}{n}} \right) \leq c_2 e^{-c_3 p},$$

where the constants depend only on the eigenvalues of Σ , bound B and radius R and sub-gaussian norm K .

Proof of Lemma C.3. We start by defining the brackets of the form

$$l_\beta(x) = f(\langle x, \beta \rangle) - \epsilon \frac{\|x\|_2}{4\mathbb{E}[\|x\|_2]},$$

$$u_\beta(x) = f(\langle x, \beta \rangle) + \epsilon \frac{\|x\|_2}{4\mathbb{E}[\|x\|_2]}.$$

Observe that the size of bracket $[\ell_\beta, u_\beta]$ is $\epsilon/2$. Now let T_Δ be a Δ -net over $B_p(R)$ where we use $\Delta = \epsilon/(4L\mathbb{E}[\|x\|_2])$. Then $\forall \beta \in B_p(R)$, $\exists \beta' \in T_\Delta$ such that $f(\langle \cdot, \beta \rangle)$ belongs to the bracket $[\ell_{\beta'}, u_{\beta'}]$. This can be seen by writing out the Lipschitz property of the function f . That is,

$$\begin{aligned} \|f(\langle x, \beta \rangle) - f(\langle x, \beta' \rangle)\|_2 &\leq L\|x\|_2\|\beta - \beta'\|_2, \\ &\leq \Delta L\|x\|_2, \end{aligned}$$

where the first inequality follows from Cauchy-Schwartz. Therefore, we conclude that

$$\mathcal{N}_{[]}(\epsilon/2, \mathcal{F}, L_1) \leq |T_\Delta|$$

for the function class $\mathcal{F} = \{f(\langle \cdot, \beta \rangle) : \beta \in B_p(R)\}$. We further have $\forall \beta \in B_p(R)$, $\exists \beta' \in T_\Delta$ such that

$$\begin{aligned} \frac{1}{n} \sum_{i=1}^n f(\langle x_i, \beta \rangle) - \mathbb{E}[f(\langle x, \beta \rangle)] &\leq \frac{1}{n} \sum_{i=1}^n u_{\beta'}(x_i) - \mathbb{E}[u_{\beta'}(x)] + \frac{\epsilon}{2}, \\ \frac{1}{n} \sum_{i=1}^n f(\langle x_i, \beta \rangle) - \mathbb{E}[f(\langle x, \beta \rangle)] &\geq \frac{1}{n} \sum_{i=1}^n l_{\beta'}(x_i) - \mathbb{E}[l_{\beta'}(x)] - \frac{\epsilon}{2}. \end{aligned}$$

Using the above inequalities, we have, $\forall \beta \in B_p(R)$, $\exists \beta' \in T_\Delta$

$$\begin{aligned} &\left\{ \left[\frac{1}{n} \sum_{i=1}^n u_{\beta'}(x_i) - \mathbb{E}[u_{\beta'}(x)] \right] > \epsilon/2 \right\} \cup \left\{ \left[-\frac{1}{n} \sum_{i=1}^n l_{\beta'}(x_i) + \mathbb{E}[l_{\beta'}(x)] \right] > \epsilon/2 \right\} \supset \\ &\left\{ \left| \frac{1}{n} \sum_{i=1}^n f(\langle x_i, \beta \rangle) - \mathbb{E}[f(\langle x, \beta \rangle)] \right| > \epsilon \right\}. \end{aligned}$$

By the union bound, we obtain

$$\begin{aligned} &\mathbb{P} \left(\max_{\beta \in T_\Delta} \left[\frac{1}{n} \sum_{i=1}^n u_\beta(x_i) - \mathbb{E}[u_\beta(x)] \right] > \epsilon/2 \right) + \mathbb{P} \left(\max_{\beta \in T_\Delta} \left[-\frac{1}{n} \sum_{i=1}^n l_\beta(x_i) + \mathbb{E}[l_\beta(x)] \right] > \epsilon/2 \right) \geq \\ &\mathbb{P} \left(\sup_{\beta \in B_p(R)} \left| \frac{1}{n} \sum_{i=1}^n f(\langle x_i, \beta \rangle) - \mathbb{E}[f(\langle x, \beta \rangle)] \right| > \epsilon \right). \end{aligned} \tag{C.3}$$

In order to complete the proof, we need concentration inequalities for u_β and l_β . We state the following lemma.

Lemma C.4. *There exists a constant C depending on the eigenvalues of Σ and B such that, for each $\beta \in B_n(R)$ and for some $0 < \epsilon < 1$, we have*

$$\begin{aligned}\mathbb{P}\left(\left|\frac{1}{n}\sum_{i=1}^n u_\beta(x_i) - \mathbb{E}[u_\beta(x)]\right| > \epsilon/2\right) &\leq 2e^{-Cn\epsilon^2}, \\ \mathbb{P}\left(\left|\frac{1}{n}\sum_{i=1}^n l_\beta(x_i) - \mathbb{E}[l_\beta(x)]\right| > \epsilon/2\right) &\leq 2e^{-Cn\epsilon^2},\end{aligned}$$

where

$$C = \frac{c}{\left(B + \frac{\sqrt{2}K}{4\mu/\sqrt{p}}\right)^2}$$

for an absolute constant c .

Remark 3. Note that $\mu = \mathbb{E}[\|x\|_2] = \mathcal{O}(\sqrt{p})$ and dividing it by \sqrt{p} would give a constant independent of n and p .

Proof of Lemma C.4. By the relation between sub-gaussian and sub-exponential norms, we have

$$\begin{aligned}\| \|x\|_2 \|_{\psi_2}^2 &\leq \| \|x\|_2^2 \|_{\psi_1} \leq \sum_{i=1}^p \|x_i^2\|_{\psi_1}, \\ &\leq 2 \sum_{i=1}^p \|x_i\|_{\psi_2}^2, \\ &\leq 2K^2 p.\end{aligned}\tag{C.4}$$

Therefore $\|x\|_2 - \mathbb{E}[\|x\|_2]$ is a centered sub-gaussian random variable with sub-gaussian norm bounded above by $2K\sqrt{2p}$. We have,

$$\mathbb{E}[\|x\|_2] = \mu.$$

Note that μ is actually of order \sqrt{p} . Assuming that the left hand side of the above equality is equal to $\sqrt{p}K'$ for some constant $K' > 0$, we can conclude that the random variable $u_\beta(x) = f(\langle x, \beta \rangle) + \epsilon \frac{\|x\|_2}{4\mathbb{E}[\|x\|_2]}$ is also sub-gaussian with

$$\begin{aligned}\|u_\beta(x)\|_{\psi_2} &\leq B + \frac{\epsilon}{4\mathbb{E}[\|x\|_2]} \| \|x\|_2 \|_{\psi_2} \\ &\leq B + \frac{\epsilon}{4\sqrt{p}K'} K\sqrt{2p} \\ &\leq B + C'\end{aligned}$$

where $C' = \sqrt{2}K/4K'$ is a constant and we also assumed $\epsilon < 1$. Now, define the function

$$g_\beta(x) = u_\beta(x) - \mathbb{E}[u_\beta(x)].$$

Note that $g_\beta(x)$ is a centered sub-gaussian random variable with sub-gaussian norm

$$\|g_\beta(x)\|_{\psi_2} \leq 2B + 2C'.$$

Then, by the Hoeffding-type inequality for the sub-gaussian random variables, we obtain

$$\mathbb{P}\left(\left|\frac{1}{n}\sum_{i=1}^n g_\beta(x_i)\right| > \epsilon/2\right) \leq 2e^{-c n \epsilon^2 / (B+C')^2}$$

where c is an absolute constant. The same argument follows for $l_\beta(x)$. \square

Using the above lemma with the union bound over the set T_Δ , we can write

$$\mathbb{P}\left(\sup_{\beta \in B_n(R)} \left|\frac{1}{n}\sum_{i=1}^n f(\langle x_i, \beta \rangle) - \mathbb{E}[f(\langle x, \beta \rangle)]\right| > \epsilon\right) \leq 4|T_\Delta|e^{-C n \epsilon^2}.$$

Since we can also write, by Lemma G.5

$$\begin{aligned} |T_\Delta| &\leq \left(\frac{R\sqrt{p}}{\Delta}\right)^p \leq \left(\frac{4RL\mathbb{E}[\|x\|_2]\sqrt{p}}{\epsilon}\right)^p, \\ &\leq \left(\frac{4\sqrt{2}RLKp}{\epsilon}\right)^p, \end{aligned}$$

and we observe that, for the constant $c' = 4\sqrt{2}RLK$,

$$\begin{aligned} \mathbb{P}\left(\sup_{\beta \in B_n(R)} \left|\frac{1}{n}\sum_{i=1}^n f(\langle x_i, \beta \rangle) - \mathbb{E}[f(\langle x_1, \beta \rangle)]\right| > \epsilon\right) &\leq 4\left(\frac{4\sqrt{2}RLKp}{\epsilon}\right)^p e^{-C n \epsilon^2}, \\ &= 4 \exp\{p \log(c'p/\epsilon) - C n \epsilon^2\}. \end{aligned}$$

We will obtain an exponential decay of order p on the right hand side. For some constant h depending on n and p , if we choose $\epsilon = hp$, we need

$$h^2 \geq \frac{1}{Cnp} \log(c'/h).$$

By the Lemma G.6, choosing $h^2 = \log(2c'^2Cnp)/(2Cnp)$, we satisfy the above requirement. Note that for n large enough, the condition of the lemma is easily satisfied. Hence, for

$$\epsilon^2 = \frac{p \log(2c'^2Cnp)}{2Cn} = \mathcal{O}\left(\frac{p \log(n)}{n}\right),$$

we obtain that there exists constants c_1, c_2, c_3 such that

$$\mathbb{P}\left(\sup_{\beta \in B_n(R)} \left|\frac{1}{n}\sum_{i=1}^n f(\langle x_i, \beta \rangle) - \mathbb{E}[f(\langle x_1, \beta \rangle)]\right| > c_1 \sqrt{\frac{p \log(n)}{n}}\right) \leq c_2 e^{-c_3 p},$$

where

$$\begin{aligned} c_1 &= \frac{3 \left(B + \frac{\sqrt{2}K}{4\sqrt{\text{Tr}(\Sigma)/p-16K^2}} \right)^2}{2c}, \\ c_2 &= 4, \\ c_3 &= \frac{1}{2} \log(7) \leq \frac{1}{2} \log(\log(64R^2L^2K^2C) + 6\log(p)). \end{aligned}$$

when $p > e$ and $64R^2L^2K^2C > e$.

□

In the following, we state the concentration results on the unbounded functions of the form

$$x \rightarrow f(\langle x, \beta \rangle) \langle x, v \rangle^2.$$

Functions of this type form the summands of the Hessian matrix in GLMs.

Lemma C.5. *Let x_i , for $i = 1, \dots, n$, be i.i.d sub-gaussian random variables with mean 0, covariance matrix Σ and sub-gaussian norm K . Also let $f : \mathbb{R} \rightarrow \mathbb{R}$ be a uniformly bounded function such that for some $B > 0$, we have $\|f\|_\infty < B$ and f is Lipschitz continuous with constant L . Further, let $v \in \mathbb{R}^p$ such that $\|v\|_2 = 1$. Then, for n, p sufficiently large satisfying*

$$n^{0.2} / \log(n) \gtrsim p$$

, there exists absolute constants c_1, c_2 depending on L, B, R and the eigenvalues of Σ such that, we have

$$\mathbb{P} \left(\sup_{\beta \in B_p(R)} \left| \frac{1}{n} \sum_{i=1}^n f(\langle x_i, \beta \rangle) \langle x_i, v \rangle^2 - \mathbb{E}[f(\langle x, \beta \rangle) \langle x, v \rangle^2] \right| > c_1 \sqrt{\frac{p}{n^{0.2} \log(n)}} \right) \leq c_2 e^{-p}.$$

Proof of Lemma C.5. We define the brackets of the form

$$\begin{aligned} l_\beta(x) &= f(\langle x, \beta \rangle) \langle x, v \rangle^2 - \epsilon \frac{\|x\|_2^3}{4\mathbb{E}[\|x\|_2^3]}, \\ u_\beta(x) &= f(\langle x, \beta \rangle) \langle x, v \rangle^2 + \epsilon \frac{\|x\|_2^3}{4\mathbb{E}[\|x\|_2^3]}, \end{aligned} \tag{C.5}$$

and we observe that the bracket $[\ell_\beta, u_\beta]$ has size $\epsilon/2$ in L_1 , that is,

$$\mathbb{E} [|u_\beta(x) - l_\beta(x)|] = \epsilon/2.$$

Next, for the following constant

$$\Delta = \frac{\epsilon}{4L\mathbb{E}[\|x\|_2^3]},$$

we define a Δ -net over $B_p(R)$ and call it \mathcal{T}_Δ . Then, $\forall \beta \in B_p(R)$, $\exists \beta' \in \mathcal{T}_\Delta$ such that $f(\langle \cdot, \beta \rangle) \langle \cdot, v \rangle^2$ belongs to the bracket $[\ell_{\beta'}, u_{\beta'}]$. This can be seen by writing the Lipschitz continuity of the function f , i.e.,

$$\begin{aligned} |f(\langle x, \beta \rangle) \langle x, v \rangle^2 - f(\langle x, \beta' \rangle) \langle x, v \rangle^2| &= \langle x, v \rangle^2 |\{f(\langle x, \beta \rangle) - f(\langle x, \beta' \rangle)\}|, \\ &\leq L \|x\|_2^2 \|v\|_2^2 |\langle x, \beta - \beta' \rangle|, \\ &\leq L \|x\|_2^3 \|\beta - \beta'\|_2, \\ &\leq \Delta L \|x\|_2^3, \end{aligned}$$

where we used Cauchy-Schwartz to obtain the above inequalities. Hence, we may conclude that for the bracketing functions in C.5, bracketing number of the function class

$$\mathcal{F} = \{f(\langle \cdot, \beta \rangle) \langle \cdot, v \rangle^2 : \beta \in B_p(R)\}$$

is bounded above by the covering number of the ball of radius R for the given scale $\Delta = \epsilon / (4L\mathbb{E}[\|x\|_2^3])$, i.e.,

$$\mathcal{N}_{[]}(\epsilon/2, \mathcal{F}, L_1) \leq |\mathcal{T}_\Delta|.$$

Next, we will upper bound the target probability using the bracketing functions u_β, l_β . We have $\forall \beta \in B_p(R)$, $\exists \beta' \in \mathcal{T}_\Delta$ such that

$$\begin{aligned} \frac{1}{n} \sum_{i=1}^n f(\langle x_i, \beta \rangle) \langle x_i, v \rangle^2 - \mathbb{E}[f(\langle x, \beta \rangle) \langle x, v \rangle^2] &\leq \frac{1}{n} \sum_{i=1}^n u_{\beta'}(x_i) - \mathbb{E}[u_{\beta'}(x)] + \frac{\epsilon}{2}, \\ \frac{1}{n} \sum_{i=1}^n f(\langle x_i, \beta \rangle) \langle x_i, v \rangle^2 - \mathbb{E}[f(\langle x, \beta \rangle) \langle x, v \rangle^2] &\geq \frac{1}{n} \sum_{i=1}^n l_{\beta'}(x_i) - \mathbb{E}[l_{\beta'}(x)] - \frac{\epsilon}{2}. \end{aligned}$$

Using the above inequalities, $\forall \beta, \exists \beta'$, we can write

$$\begin{aligned} &\left\{ \left[\frac{1}{n} \sum_{i=1}^n u_{\beta'}(x_i) - \mathbb{E}[u_{\beta'}(x)] \right] > \epsilon/2 \right\} \cup \left\{ \left[-\frac{1}{n} \sum_{i=1}^n l_{\beta'}(x_i) + \mathbb{E}[l_{\beta'}(x)] \right] > \epsilon/2 \right\} \supset \\ &\left\{ \left| \frac{1}{n} \sum_{i=1}^n f(\langle x_i, \beta \rangle) \langle x_i, v \rangle^2 - \mathbb{E}[f(\langle x, \beta \rangle) \langle x, v \rangle^2] \right| > \epsilon \right\}. \end{aligned}$$

Hence, by the union bound, we obtain

$$\begin{aligned} &\mathbb{P} \left(\max_{\beta \in \mathcal{T}_\Delta} \left[\frac{1}{n} \sum_{i=1}^n u_\beta(x_i) - \mathbb{E}[u_\beta(x)] \right] > \epsilon/2 \right) + \mathbb{P} \left(\max_{\beta \in \mathcal{T}_\Delta} \left[-\frac{1}{n} \sum_{i=1}^n l_\beta(x_i) + \mathbb{E}[l_\beta(x)] \right] > \epsilon/2 \right) \geq \\ &\mathbb{P} \left(\sup_{\beta \in B_p(R)} \left| \frac{1}{n} \sum_{i=1}^n f(\langle x_i, \beta \rangle) \langle x_i, v \rangle^2 - \mathbb{E}[f(\langle x, \beta \rangle) \langle x, v \rangle^2] \right| > \epsilon \right). \end{aligned} \tag{C.6}$$

In order to complete the proof, we need one-sided concentration inequalities for u_β and l_β . Handling these functions is somewhat tedious since $\|x\|_2^3$ terms do not concentrate nicely. We state the following lemma.

Lemma C.6. For a given $\alpha, \epsilon > 0$, and n sufficiently large such that, $\nu(n^\alpha, p, \epsilon, B, K, \Sigma) < \epsilon/4$ where

$$\begin{aligned} \nu(n^\alpha, p, \epsilon, B, K, \Sigma) = & \left(n^\alpha + \frac{6B^2K^2p}{c} \right) \exp \left(-c \frac{n^\alpha}{6B^2K^2p} \right) \\ & + \left\{ n^\alpha + \frac{4K^2p}{c\text{Tr}(\Sigma)} n^{\alpha/3} \epsilon^{2/3} + \frac{6K^4p^2}{c^2\text{Tr}(\Sigma)^2} \epsilon^{4/3} n^{-\alpha/3} \right\} \exp \left(-c \frac{\text{Tr}(\Sigma)(n^\alpha/\epsilon)^{2/3}}{2K^2p} \right). \end{aligned}$$

Then, there exists constants c', c'', c''' depending on the eigenvalues of Σ , B and K such that $\forall \beta$, we have,

$$\begin{aligned} \mathbb{P} \left(\frac{1}{n} \sum_{i=1}^n u_\beta(x_i) - \mathbb{E}[u_\beta(x)] > \epsilon/2 \right) & \leq \exp(-c'n^\alpha/p) + \exp(-c''n^{2\alpha/3}\epsilon^{-2/3}) + \exp(-c'''n^{1-2\alpha}\epsilon^2), \\ \mathbb{P} \left(-\frac{1}{n} \sum_{i=1}^n l_\beta(x_i) + \mathbb{E}[l_\beta(x)] > \epsilon/2 \right) & \leq \exp(-c'n^\alpha/p) + \exp(-c''n^{2\alpha/3}\epsilon^{-2/3}) + \exp(-c'''n^{1-2\alpha}\epsilon^2). \end{aligned}$$

Proof of Lemma C.6. For the sake of simplicity, we define the functions

$$\begin{aligned} \tilde{u}_\beta(w) &= u_\beta(w) - \mathbb{E}[u_\beta(x)], \\ \tilde{l}_\beta(w) &= l_\beta(w) - \mathbb{E}[l_\beta(x)]. \end{aligned}$$

We will show the concentration of the upper bracket, \tilde{u} . Proof for the lower bracket \tilde{l} , will follow from similar steps. We write,

$$\begin{aligned} \mathbb{P} \left(\frac{1}{n} \sum_{i=1}^n \tilde{u}_\beta(x_i) > \epsilon/2 \right) & \leq \mathbb{P} \left(\frac{1}{n} \sum_{i=1}^n \tilde{u}_\beta(x_i) > \epsilon/2, \max_{1 \leq i \leq n} |\tilde{u}_\beta(x_i)| < n^\alpha \right) \\ & \quad + \mathbb{P} \left(\max_{1 \leq i \leq n} |\tilde{u}_\beta(x_i)| \geq n^\alpha \right). \end{aligned} \tag{C.7}$$

We need to bound the right hand side of the above equation. For the second term, since $\tilde{u}_\beta(x_i)$'s are i.i.d. centered random variables, we have

$$\begin{aligned} \mathbb{P} \left(\max_{1 \leq i \leq n} |\tilde{u}_\beta(x_i)| \geq n^\alpha \right) &= 1 - \mathbb{P} \left(\max_{1 \leq i \leq n} |\tilde{u}_\beta(x_i)| < n^\alpha \right), \\ &= 1 - \mathbb{P}(|\tilde{u}_\beta(x)| < n^\alpha)^n, \\ &= 1 - (1 - \mathbb{P}(|\tilde{u}_\beta(x)| \geq n^\alpha))^n, \\ &\leq n\mathbb{P}(|\tilde{u}_\beta(x)| \geq n^\alpha). \end{aligned}$$

Also, note that

$$\begin{aligned} |\tilde{u}_\beta(x)| &\leq B\|x\|_2^2 + \epsilon \frac{\|x\|_2^3}{4\mathbb{E}[\|x\|_2^3]} + \mathbb{E}[u_\beta(x)], \\ &\leq B\|x\|_2^2 + \epsilon \frac{\|x\|_2^3}{4\mathbb{E}[\|x\|_2^3]} + B\text{Tr}(\Sigma) + \epsilon/4. \end{aligned}$$

Therefore, if $t > 3B\text{Tr}(\mathbf{\Sigma})$ and for ϵ small, we can write

$$\{|\tilde{u}_\beta(x)| > t\} \subset \{B\|x\|_2^2 > t/3\} \cup \left\{ \epsilon \frac{\|x\|_2^3}{4\mathbb{E}[\|x\|_2^3]} > t/3 \right\}. \quad (\text{C.8})$$

Since x is a sub-gaussian random variable with $\|x\|_{\psi_2} = K$, we have

$$K = \sup_{w \in S^{p-1}} \|\langle w, x \rangle\|_{\psi_2} = \|x\|_{\psi_2}.$$

Using this and the relation between sub-gaussian and sub-exponential norms as in Eq. (C.4), we have $\|\|x\|_2\|_{\psi_2}^2 \leq 2K^2p$. This provides the following tail bound for $\|x\|_2$,

$$\mathbb{P}(\|x\|_2 > s) \leq \exp\left(-\frac{cs^2}{2pK^2}\right), \quad (\text{C.9})$$

where c is an absolute constant. Using the above tail bound, we can write,

$$\mathbb{P}\left(\|x\|_2^2 > \frac{1}{3B}t\right) \leq 2\exp\left(-c\frac{t}{6BK^2p}\right).$$

For the next term in Eq. (C.8), we need a lower bound for $\mathbb{E}[\|x\|_2^3]$. We use a modified version of the Hölder's inequality and obtain

$$\mathbb{E}[\|x\|_2^3] \geq \mathbb{E}[\|x\|_2^2]^{3/2} = \text{Tr}(\mathbf{\Sigma})^{3/2}.$$

Using the above inequality, we can write

$$\begin{aligned} \mathbb{P}\left(\epsilon \frac{\|x\|_2^3}{4\mathbb{E}[\|x\|_2^3]} > t/3\right) &\leq \mathbb{P}\left(\|x\|_2^3 > \frac{4}{3\epsilon}\text{Tr}(\mathbf{\Sigma})^{3/2}t\right), \\ &= \mathbb{P}\left(\|x\|_2 > \left(\frac{4t}{3\epsilon}\right)^{1/3} \text{Tr}(\mathbf{\Sigma})^{1/2}\right), \\ &\leq \exp\left(-c\frac{\text{Tr}(\mathbf{\Sigma})(t/\epsilon)^{2/3}}{2K^2p}\right). \end{aligned}$$

where c is the same absolute constant as in Eq. (C.9).

Now for $\alpha > 0$ such that $t = n^\alpha > 3B\text{Tr}(\mathbf{\Sigma})$ (we will justify this assumption for a particular choice of α later), we combine the above results,

$$\mathbb{P}(|\tilde{u}_\beta(x)| > t) \leq 2\exp\left(-c\frac{t}{6BK^2p}\right) + 2\exp\left(-c\frac{\text{Tr}(\mathbf{\Sigma})(t/\epsilon)^{2/3}}{2K^2p}\right). \quad (\text{C.10})$$

Next, we focus on the first term in Eq.(C.7). Let $\mu = \mathbb{E}[\tilde{u}_\beta(x)\mathbb{1}_{\{|\tilde{u}_\beta(x)| < n^\alpha\}}]$, and write

$$\begin{aligned} \mathbb{P}\left(\frac{1}{n}\sum_{i=1}^n \tilde{u}_\beta(x_i) > \epsilon/2 ; \max_{1 \leq i \leq n} |\tilde{u}_\beta(x_i)| < n^\alpha\right) &\leq \mathbb{P}\left(\frac{1}{n}\sum_{i=1}^n \tilde{u}_\beta(x_i)\mathbb{1}_{\{|\tilde{u}_\beta(x_i)| < n^\alpha\}} > \epsilon/2\right), \\ &= \mathbb{P}\left(\frac{1}{n}\sum_{i=1}^n \tilde{u}_\beta(x_i)\mathbb{1}_{\{|\tilde{u}_\beta(x_i)| < n^\alpha\}} - \mu > \epsilon/2 - \mu\right), \\ &\leq \exp\left\{-\frac{n^{1-2\alpha}}{2}\left(\frac{\epsilon}{2} - \mu\right)^2\right\}, \end{aligned}$$

where we used the Hoeffding's concentration inequality for the bounded random variables. Further, note that

$$0 = \mathbb{E}[\tilde{u}_\beta(x)] = \mu + \mathbb{E}\left[\tilde{u}_\beta(x)\mathbb{1}_{\{|\tilde{u}_\beta(x)| > n^\alpha\}}\right].$$

By Lemma G.2, we can write

$$|\mu| = \left|\mathbb{E}\left[\tilde{u}_\beta(x)\mathbb{1}_{\{|\tilde{u}_\beta(x)| > n^\alpha\}}\right]\right| \leq n^\alpha \mathbb{P}(|\tilde{u}_\beta(x)| > n^\alpha) + \int_{n^\alpha}^{\infty} \mathbb{P}(|\tilde{u}_\beta(x)| > t) dt.$$

The first term on the right hand side can be easily bounded by using Eq.(C.10), i.e.,

$$n^\alpha \mathbb{P}(|\tilde{u}_\beta(x)| > n^\alpha) \leq n^\alpha \exp\left(-c \frac{n^\alpha}{6BK^2p}\right) + n^\alpha \exp\left(-c \frac{\text{Tr}(\mathbf{\Sigma})(n^\alpha/\epsilon)^{2/3}}{2K^2p}\right).$$

For the second term, using Eq.(C.10) once again, we obtain

$$\begin{aligned} \int_{n^\alpha}^{\infty} \mathbb{P}(|\tilde{u}_\beta(x)| > t) dt &\leq \int_{n^\alpha}^{\infty} \exp(-c't/p) dt + \int_{n^\alpha}^{\infty} \exp\left(-c \frac{\text{Tr}(\mathbf{\Sigma})(t/\epsilon)^{2/3}}{4K^2p}\right) dt, \\ &= \frac{6BK^2p}{c} \exp\left(-c \frac{n^\alpha}{6BK^2p}\right) + \int_{n^\alpha}^{\infty} \exp\left(-c \frac{\text{Tr}(\mathbf{\Sigma})(t/\epsilon)^{2/3}}{2K^2p}\right) dt. \end{aligned}$$

Next, we apply Lemma G.3 to bound the second term on the right hand side. That is, we have

$$\int_{n^\alpha}^{\infty} \exp\left(-c''(t/\epsilon)^{2/3}\right) dt \leq \left\{ \frac{3K^2p}{c\text{Tr}(\mathbf{\Sigma})} n^{\alpha/3} \epsilon^{2/3} + \frac{6K^4p^2}{c^2\text{Tr}(\mathbf{\Sigma})^2} \epsilon^{4/3} n^{-\alpha/3} \right\} \exp\left(-c \frac{\text{Tr}(\mathbf{\Sigma})(n^\alpha/\epsilon)^{2/3}}{2K^2p}\right).$$

Combining the above results, we can write

$$\begin{aligned} |\mu| &\leq \left(n^\alpha + \frac{6B^2K^2p}{c}\right) \exp\left(-c \frac{n^\alpha}{6B^2K^2p}\right) \\ &\quad + \left\{ n^\alpha + \frac{4K^2p}{c\text{Tr}(\mathbf{\Sigma})} n^{\alpha/3} \epsilon^{2/3} + \frac{6K^4p^2}{c^2\text{Tr}(\mathbf{\Sigma})^2} \epsilon^{4/3} n^{-\alpha/3} \right\} \exp\left(-c \frac{\text{Tr}(\mathbf{\Sigma})(n^\alpha/\epsilon)^{2/3}}{2K^2p}\right), \\ &=: \nu(n^\alpha, p, \epsilon, B, K, \mathbf{\Sigma}). \end{aligned}$$

Notice that, the upper bound on $|\mu|$, namely $\nu(n^\alpha, p, \epsilon, B, K, \mathbf{\Sigma})$, is close to 0 when n is large. This is because exponentially decaying functions will dominate the coefficients. We assume that n is sufficiently large that the upper bound for $|\mu|$ is less than $\epsilon/4$. In particular, we will choose $\alpha = 0.4$ later in the proof.

Applying this bounds on Eq.(C.7), we obtain

$$\begin{aligned} \mathbb{P}\left(\frac{1}{n} \sum_{i=1}^n \tilde{u}_\beta(x_i) > \epsilon/2\right) &\leq \exp\left(-c \frac{n^\alpha}{6B^2K^2p}\right) + \exp\left(-c \frac{\text{Tr}(\mathbf{\Sigma})(n^\alpha/\epsilon)^{2/3}}{2K^2p}\right) + \exp\left(-\frac{n^{1-2\alpha}}{8} \epsilon^2\right), \\ &= \exp(-c'n^\alpha/p) + \exp(-c''n^{2\alpha/3}\epsilon^{-2/3}) + \exp(-c'''n^{1-2\alpha}\epsilon^2), \end{aligned}$$

where

$$\begin{aligned} c' &= \frac{c}{6B^2K^2}, \\ c'' &= \frac{c\text{Tr}(\Sigma)/p}{2K^2} \geq \frac{c\sigma^2}{2K^2}, \\ c''' &= \frac{1}{8}. \end{aligned}$$

Hence, the proof is completed for the upper bracket.

The proof for the lower brackets $l_\beta(x)$ follows from exactly the same steps and omitted here. \square

Applying the above lemma on Eq.(C.6), for $\alpha > 0$, we obtain

$$\begin{aligned} \mathbb{P} \left(\sup_{\beta \in B_n(R)} \left| \frac{1}{n} \sum_{i=1}^n f(\langle x_i, \beta \rangle) \langle x_i, v \rangle^2 - \mathbb{E}[f(\langle x, \beta \rangle) \langle x, v \rangle^2] \right| > \epsilon \right) &\leq 2|T_\Delta| \exp(-c'n^\alpha/p) \\ &\quad + 2|T_\Delta| \exp(-c''n^{2\alpha/3}\epsilon^{-2/3}) \\ &\quad + 2|T_\Delta| \exp(-c'''n^{1-2\alpha}\epsilon^2). \end{aligned} \quad (\text{C.11})$$

Observe that we can write, by Lemma G.5

$$|T_\Delta| \leq \left(\frac{R\sqrt{p}}{\Delta} \right)^p = \left(\frac{4\sqrt{p}RL\mathbb{E}[\|x\|_2^3]}{\epsilon} \right)^p.$$

Also, recall that $\|x\|_2$ was a sub-gaussian random variable with $\|\|x\|_2\|_{\psi_2} \leq K\sqrt{2p}$. Using the definition of sub-gaussian norm, we have

$$\begin{aligned} \frac{1}{\sqrt{3}} \mathbb{E}[\|x\|_2^3]^{1/3} &\leq \|\|x\|_2\|_{\psi_2} \leq \sqrt{2p}K, \\ \implies \mathbb{E}[\|x\|_2^3] &\leq 15K^3p^{3/2}. \end{aligned}$$

Therefore, we have $\mathbb{E}[\|x\|_2^3] = \mathcal{O}(p^{3/2})$ (recall that we had a lower bound of the same order). We define a constant K' , and as ϵ is small, we have

$$\begin{aligned} |T_\Delta| &\leq \left(\frac{60RLK^3p^2}{\epsilon} \right)^p \\ &= \left(\frac{K'p^2}{\epsilon} \right)^p, \end{aligned}$$

where we let $K' = 60RLK^3$. We will show that each term on the right hand side of Eq.(C.11) decays exponentially with a rate of order p . For the first term, for $s > 0$, we write

$$\begin{aligned} |T_\Delta| \exp(-c'n^\alpha/p) &= \exp(-c'n^\alpha/p + p \log(K') + 2p \log(p) + p \log(\epsilon^{-1})), \\ &\leq \exp(-c'n^\alpha/p + 2p \log(K'p/\epsilon)). \end{aligned} \quad (\text{C.12})$$

Similarly for the second and third terms, we write

$$\begin{aligned} |T_\Delta| \exp\left(-c'' n^{2\alpha/3} \epsilon^{-2/3}\right) &\leq \exp\left(-c'' n^{2\alpha/3} \epsilon^{-2/3} + 2p \log(K' p / \epsilon)\right), \\ |T_\Delta| \exp\left(-c''' n^{1-2\alpha} \epsilon^2\right) &\leq \exp\left(-c''' n^{1-2\alpha} \epsilon^2 + 2p \log(K' p / \epsilon)\right). \end{aligned} \quad (\text{C.13})$$

We will seek values for ϵ and α to obtain an exponential decay with rate p on the right sides of Eq.(C.12,C.13). That is, we need

$$\begin{aligned} c' n^\alpha / p &\geq 2p \log(K'' p / \epsilon), \\ c'' n^{2\alpha/3} &\geq 2p \log(K'' p / \epsilon) \epsilon^{2/3}, \\ c''' n^{1-2\alpha} \epsilon^2 &\geq 2p \log(K'' p / \epsilon), \end{aligned} \quad (\text{C.14})$$

where $K'' = eK'$.

We apply Lemma G.6 for the last inequality in Eq. (C.14). That is,

$$\begin{aligned} \epsilon^2 &= \frac{p}{c''' n^{1-2\alpha}} \log(c''' K'' n^{1-2\alpha}), \\ &= \mathcal{O}\left(\frac{p}{n^{1-2\alpha}} \log(n)\right). \end{aligned}$$

where we assume that

$$n^{1-2\alpha} p > \frac{e}{c''' K''^2}.$$

The above statement holds for $\alpha < 1/2$ and n large enough, as the right hand side is an absolute constant.

In the following, we choose $\alpha = 0.4$ and use the assumption that

$$n^{0.2} / \log(n) \gtrsim p,$$

which provides $\epsilon < 1$. Note that this choice of α also justifies the assumption used to derive Eq. (C.10). One can easily check that $\alpha = 0.4$ implies that the first and the second statements in Eq. (C.14) are satisfied.

It remains to check whether $\nu(n^\alpha, p, \epsilon, B, K, \Sigma) < \epsilon/4$ (in Lemma C.6) for this particular choice of α and ϵ . It suffices to consider only the dominant terms in the definition of ν . We use the assumption on n, p and write

$$\begin{aligned} \nu(n^{0.2}, p, \epsilon, B, K, \Sigma) &\lesssim n^{0.2} \exp\left(-\frac{cn^{0.4}}{6B^2 K^2 p}\right) + n^{0.2} \exp\left(-\frac{c\text{Tr}(\Sigma)/p n^{0.8/3}}{2K^2}\right), \\ &\lesssim n^{0.2} \exp\left(-\frac{c}{6B^2 K^2} n^{0.2}\right) + n^{0.2} \exp\left(-\frac{c\sigma^2}{2K^2} n^{0.8/3}\right). \end{aligned}$$

For n sufficiently large, this quantity is always smaller than $\epsilon/4$. Hence, for some constant c_1, c_2 , we obtain

$$\mathbb{P}\left(\sup_{\beta \in B_p(R)} \left|\frac{1}{n} \sum_{i=1}^n f(\langle x_i, \beta \rangle) \langle x_i, v \rangle^2 - \mathbb{E}[f(\langle x, \beta \rangle) \langle x, v \rangle^2]\right| > c_1 \sqrt{\frac{p}{n^{0.2}} \log(n)}\right) \leq c_2 e^{-p},$$

where

$$\begin{aligned} c_1 &= 16, \\ c_2 &= 6. \end{aligned}$$

□

D Proof of main theorem

We will provide the proofs of Theorems 4.1 and 4.4 in parallel. Matrix concentration results in this section are mostly based on the covering net argument provided in [Ver10]. Similar results for matrix forms can also be obtained through different techniques such as *chaining* as well (See i.e. [DE15]). On the set \mathcal{E} , we write,

$$\begin{aligned} \hat{\beta}^t - \beta_* - \gamma \mathbf{Q}^t \nabla_{\beta} \ell(\hat{\beta}^t) &= \hat{\beta}^t - \beta_* - \gamma \mathbf{Q}^t \int_0^1 \nabla_{\beta}^2 \ell(\beta_* + \xi(\hat{\beta}^t - \beta_*)) d\xi (\hat{\beta}^t - \beta_*), \\ &= \left(I - \gamma \mathbf{Q}^t \int_0^1 \nabla_{\beta}^2 \ell(\beta_* + \xi(\hat{\beta}^t - \beta_*)) d\xi \right) (\hat{\beta}^t - \beta_*). \end{aligned}$$

Since the projection $\mathcal{P}_{B_p(R)}$ in step 3 of NewSt can only decrease the ℓ_2 distance, we obtain

$$\|\hat{\beta}^{t+1} - \beta_*\|_2 \leq \left\| I - \gamma \mathbf{Q}^t \int_0^1 \nabla_{\beta}^2 \ell(\beta_* + \xi(\hat{\beta}^t - \beta_*)) d\xi \right\|_2 \|\hat{\beta}^t - \beta_*\|_2. \quad (\text{D.1})$$

The governing term (with $\gamma = 1$) that determines the convergence rate can be bounded as

$$\left\| I - \mathbf{Q}^t \int_0^1 \nabla_{\beta}^2 \ell(\beta_* + \xi(\hat{\beta}^t - \beta_*)) d\xi \right\|_2 \leq \left\| [\mathbf{Q}^t]^{-1} - \int_0^1 \nabla_{\beta}^2 \ell(\beta_* + \xi(\hat{\beta}^t - \beta_*)) d\xi \right\|_2 \|\mathbf{Q}^t\|_2.$$

We define the following,

$$\mathfrak{E}(\beta) = \mathbb{E} \left[\phi^{(2)}(\langle x, \beta \rangle) \right] \Sigma + \mathbb{E} \left[\phi^{(4)}(\langle x, \beta \rangle) \right] \Sigma \beta \beta^T \Sigma$$

Note that for a function f , $\mathbb{E}[f(\langle x, \beta \rangle)] = h(\beta)$ is a function of β . With a slight abuse of notation, we write $\mathbb{E}[f(\langle x, \hat{\beta} \rangle)] = h(\hat{\beta})$ as a random variable. We have

$$\begin{aligned} \left\| [\mathbf{Q}^t]^{-1} - \int_0^1 \nabla_{\beta}^2 \ell(\beta_* + \xi(\hat{\beta}^t - \beta_*)) d\xi \right\|_2 &\leq \left\| [\mathbf{Q}^t]^{-1} - \mathfrak{E}(\hat{\beta}^t) \right\|_2 \\ &\quad + \left\| \mathbb{E}[xx^T \phi^{(2)}(\langle x, \hat{\beta}^t \rangle)] - \mathfrak{E}(\hat{\beta}^t) \right\|_2 \\ &\quad + \left\| \int_0^1 \nabla_{\beta}^2 \ell(\beta_* + \xi(\hat{\beta}^t - \beta_*)) d\xi - \mathbb{E} \left[xx^T \int_0^1 \phi^{(2)}(\langle x, \beta_* + \xi(\hat{\beta}^t - \beta_*) \rangle) d\xi \right] \right\|_2 \\ &\quad + \left\| \mathbb{E}[xx^T \phi^{(2)}(\langle x, \hat{\beta}^t \rangle)] - \mathbb{E} \left[xx^T \int_0^1 \phi^{(2)}(\langle x, \beta_* + \xi(\hat{\beta}^t - \beta_*) \rangle) d\xi \right] \right\|_2. \end{aligned}$$

For the first term on the right hand side, we state the following lemma.

Lemma D.1. *In the case of sub-gaussian covariates, there exist constants C_1, C_2 such that, with probability at least $1 - C_1/p^2$,*

$$\left\| [\mathbf{Q}^t]^{-1} - \mathfrak{E}(\hat{\beta}^t) \right\|_2 \leq C_2 \sqrt{\frac{p}{\min\{|S|p/\log(p), n/\log(n)\}}}.$$

Similarly, when the covariates are sampled from a distribution with bounded support, there exists constants C'_1, C'_2, C'_3 such that, with probability $1 - C'_1 e^{-C'_2 p}$,

$$\left\| [\mathbf{Q}^t]^{-1} - \mathfrak{E}(\hat{\beta}^t) \right\|_2 \leq C'_3 \sqrt{\frac{p}{\min\{|S|, n/\log(n)\}}}.$$

where the constants depend on K, B and the radius R .

Proof of Lemma D.1. In the following, we will only provide the proof for the bounded support case. The proof for the sub-gaussian covariates can be carried by replacing Lemmas B.3 and B.4 with Lemmas B.1 and B.2.

Using a uniform bound on the feasible set, we write

$$\begin{aligned} & \left\| [\mathbf{Q}^t]^{-1} - \mathfrak{E}(\hat{\beta}^t) \right\|_2 \\ & \leq \sup_{\beta \in B_p(R)} \left\| \hat{\mu}_2(\beta) \zeta_r(\hat{\Sigma}_S) + \hat{\mu}_4(\beta) \zeta_r(\hat{\Sigma}_S) \beta (\zeta_r(\hat{\Sigma}_S) \beta)^T - \mathbb{E}[\phi^{(2)}(\langle x, \beta \rangle)] \Sigma - \mathbb{E}[\phi^{(4)}(\langle x, \beta \rangle)] \Sigma \beta \beta^T \Sigma \right\|_2. \end{aligned}$$

We will find an upper bound for the quantity inside the supremum. By denoting the expectations of $\hat{\mu}_2(\beta)$ and $\hat{\mu}_4(\beta)$, with $\mu_2(\beta)$ and $\mu_4(\beta)$ respectively, we write

$$\begin{aligned} & \left\| \hat{\mu}_2(\beta) \zeta_r(\hat{\Sigma}_S) + \hat{\mu}_4(\beta) \zeta_r(\hat{\Sigma}_S) \beta (\zeta_r(\hat{\Sigma}_S) \beta)^T - \mathbb{E}[\phi^{(2)}(\langle x, \beta \rangle)] \Sigma - \mathbb{E}[\phi^{(4)}(\langle x, \beta \rangle)] \Sigma \beta (\Sigma \beta)^T \right\|_2 \\ & \leq \left\| \hat{\mu}_2(\beta) \zeta_r(\hat{\Sigma}_S) - \mu_2(\beta) \Sigma \right\|_2 + \left\| \hat{\mu}_4(\beta) \zeta_r(\hat{\Sigma}_S) \beta (\zeta_r(\hat{\Sigma}_S) \beta)^T - \mu_4(\beta) \Sigma \beta (\Sigma \beta)^T \right\|_2. \end{aligned}$$

For the first term on the right hand side, we have

$$\begin{aligned} \left\| \hat{\mu}_2(\beta) \zeta_r(\hat{\Sigma}_S) - \mu_2(\beta) \Sigma \right\|_2 & \leq |\hat{\mu}_2(\beta)| \left\| \zeta_r(\hat{\Sigma}_S) - \Sigma \right\|_2 + \|\Sigma\|_2 |\hat{\mu}_2(\beta) - \mu_2(\beta)|, \\ & \leq B_2 \left\| \zeta_r(\hat{\Sigma}_S) - \Sigma \right\|_2 + K |\hat{\mu}_2(\beta) - \mu_2(\beta)|. \end{aligned}$$

By the Lemmas B.3 and B.4, for some constants c_1, c_2, c_3 , we have with probability $1 - c_2 e^{-c_3 p} - 1/p^2$,

$$\begin{aligned} \sup_{\beta \in B_p(R)} \left\| \hat{\mu}_2(\beta) \zeta_r(\hat{\Sigma}_S) - \mu_2(\beta) \Sigma \right\|_2 & \leq 2B_2 K \sqrt{C} \sqrt{\frac{\log(p)}{|S|}} + c_1 K \sqrt{\frac{p \log(n)}{n}}, \\ & = \mathcal{O} \left(\sqrt{\frac{p}{\min\{p/\log(p)|S|, n/\log(n)\}}} \right). \end{aligned}$$

For the second term, we have

$$\begin{aligned}
& \left\| \hat{\mu}_4(\beta) \zeta_r(\hat{\Sigma}_S) \beta (\zeta_r(\hat{\Sigma}_S) \beta)^T - \mu_4(\beta) \Sigma \beta (\Sigma \beta)^T \right\|_2 \\
& \leq |\hat{\mu}_4(\beta)| \left\| \zeta_r(\hat{\Sigma}_S) \beta \beta^T \zeta_r(\hat{\Sigma}_S) - \Sigma \beta \beta^T \Sigma \right\|_2 + |\hat{\mu}_4(\beta) - \mu_4(\beta)| \left\| \Sigma \beta \beta^T \Sigma \right\|_2, \\
& \leq B_4 R^2 \left\{ \left\| \zeta_r(\hat{\Sigma}_S) \right\|_2 + \left\| \Sigma \right\|_2 \right\} \left\| \zeta_r(\hat{\Sigma}_S) - \Sigma \right\|_2 + R^2 \left\| \Sigma \right\|_2^2 |\hat{\mu}_4(\beta) - \mu_4(\beta)|, \\
& \leq B_4 R^2 \left\{ \left\| \zeta_r(\hat{\Sigma}_S) \right\|_2 + K \right\} \left\| \zeta_r(\hat{\Sigma}_S) - \Sigma \right\|_2 + R^2 K^2 |\hat{\mu}_4(\beta) - \mu_4(\beta)|.
\end{aligned}$$

Again, by the Lemmas B.3, B.4 and C.2, for some constants c_1, c_2, c_3 , we have with probability $1 - c_2 e^{-c_3 p} - 1/p^2$, we write

$$\begin{aligned}
B_4 R^2 \left\{ \left\| \zeta_r(\hat{\Sigma}_S) \right\|_2 + K \right\} \left\| \zeta_r(\hat{\Sigma}_S) - \Sigma \right\|_2 & \leq 2K \sqrt{C} B_4 R^2 \left\{ 2K + 3K \sqrt{C} \sqrt{\frac{\log(p)}{|S|}} \right\} \sqrt{\frac{\log(p)}{|S|}}, \\
& \leq 4K^2 \sqrt{C} B_4 R^2 \sqrt{\frac{\log(p)}{|S|}} + 6K^2 C B_4 R^2 \frac{\log(p)}{|S|}, \\
& = \mathcal{O} \left(\sqrt{\frac{\log(p)}{|S|}} \right),
\end{aligned}$$

for sufficiently large $|S|$.

Further, by Lemma C.1, for constants c_1, c_2, c_3 , we have with probability $1 - c_2 e^{-c_3 p}$,

$$\sup_{\{\beta \in B_p(R)\}} |\hat{\mu}_4(\beta) - \mu_4(\beta)| \leq c_1 \sqrt{\frac{p \log(n)}{n}} = \mathcal{O} \left(\sqrt{\frac{p \log(n)}{n}} \right).$$

Combining the above results, for sufficiently large $p, |S|$ and constants c_1, c_2 , we have with probability at least $1 - c_1/p^2$,

$$\begin{aligned}
& \sup_{\{\beta \in B_p(R)\}} \left\| \hat{\mu}_4(\beta) \zeta_r(\hat{\Sigma}_S) \beta (\zeta_r(\hat{\Sigma}_S) \beta)^T - \mu_4(\beta) \Sigma \beta (\Sigma \beta)^T \right\|_2 \\
& \leq 4K^2 \sqrt{C} \max\{B_2, B_4\} R^2 \sqrt{\frac{\log(p)}{|S|}} + 6K^2 C B_4 R^2 \frac{\log(p)}{|S|} + c_1 R^2 K^2 \sqrt{\frac{p \log(n)}{n}} \\
& = \mathcal{O} \left(\sqrt{\frac{p}{\min\{|S|p/\log(p), n/\log(n)\}}} \right).
\end{aligned}$$

Hence, for some constants C_1, C_2 , with probability $1 - C_1/p^2$, we have

$$\left\| [\mathbf{Q}^t]^{-1} - \mathfrak{E}(\hat{\beta}^t) \right\|_2 \leq C_2 \sqrt{\frac{p}{\min\{|S|p/\log(p), n/\log(n)\}}},$$

where the constants depend on $K, B = \max\{B_2, B_4\}$ and the radius R . □

Lemma D.2. *The bias term can be upper bounded by*

$$\left\| \mathbb{E}[xx^T \phi^{(2)}(\langle x, \hat{\beta}^t \rangle)] - \mathfrak{E}(\hat{\beta}^t) \right\|_2 \leq d_{\mathcal{H}_3}(x, z) + \|\Sigma\|_2 d_{\mathcal{H}_1}(x, z) + \|\Sigma\|_2^2 R^2 d_{\mathcal{H}_2}(x, z),$$

for both sub-gaussian and bounded support cases.

Proof of Lemma D.2. For a random variable $z \sim \mathbf{N}_p(0, \Sigma)$, by the triangle inequality, we write

$$\begin{aligned} & \left\| \mathbb{E}[xx^T \phi^{(2)}(\langle x, \hat{\beta}^t \rangle)] - \mathfrak{E}(\hat{\beta}^t) \right\|_2 \\ & \leq \left\| \mathbb{E}[xx^T \phi^{(2)}(\langle x, \hat{\beta}^t \rangle)] - \mathbb{E}[zz^T \phi^{(2)}(\langle z, \hat{\beta}^t \rangle)] \right\|_2 + \left\| \mathbb{E}[zz^T \phi^{(2)}(\langle z, \hat{\beta}^t \rangle)] - \mathfrak{E}(\hat{\beta}^t) \right\|_2 \end{aligned}$$

For the first term on the right hand side, we have

$$\begin{aligned} & \left\| \mathbb{E}[xx^T \phi^{(2)}(\langle x, \hat{\beta}^t \rangle)] - \mathbb{E}[zz^T \phi^{(2)}(\langle z, \hat{\beta}^t \rangle)] \right\|_2 \\ & \leq \sup_{\beta \in B_p(R)} \sup_{\|v\|_2=1} \left| \mathbb{E} \left[\langle v, x \rangle^2 \phi^{(2)}(\langle x, \beta \rangle) \right] - \mathbb{E} \left[\langle v, z \rangle^2 \phi^{(2)}(\langle z, \beta \rangle) \right] \right|, \\ & \leq d_{\mathcal{H}_3}(x, z). \end{aligned}$$

For the second term, we write

$$\begin{aligned} & \left\| \mathbb{E}[zz^T \phi^{(2)}(\langle z, \hat{\beta}^t \rangle)] - \mathfrak{E}(\hat{\beta}^t) \right\|_2 \\ & \leq \sup_{\{\beta \in B_p(R)\}} \left\| \mathbb{E}[zz^T \phi^{(2)}(\langle z, \beta \rangle)] - \mathbb{E}[\phi^{(2)}(\langle x, \beta \rangle)] \Sigma + \mathbb{E} \left[\phi^{(4)}(\langle x, \beta \rangle) \right] \Sigma \beta \beta^T \Sigma \right\|_2, \\ & \leq \sup_{\{\beta \in B_p(R)\}} \left\| \mathbb{E}[\phi^{(2)}(\langle z, \beta \rangle)] \Sigma + \mathbb{E} \left[\phi^{(4)}(\langle z, \beta \rangle) \right] \Sigma \beta \beta^T \Sigma \right. \\ & \quad \left. - \mathbb{E}[\phi^{(2)}(\langle x, \beta \rangle)] \Sigma - \mathbb{E} \left[\phi^{(4)}(\langle x, \beta \rangle) \right] \Sigma \beta \beta^T \Sigma \right\|_2, \\ & \leq \sup_{\{\beta \in B_p(R)\}} \left\| \mathbb{E}[\phi^{(2)}(\langle z, \beta \rangle)] \Sigma - \mathbb{E}[\phi^{(2)}(\langle x, \beta \rangle)] \Sigma \right\|_2, \\ & \quad + \sup_{\{\beta \in B_p(R)\}} \left\| \mathbb{E} \left[\phi^{(4)}(\langle z, \beta \rangle) \right] \Sigma \beta \beta^T \Sigma - \mathbb{E} \left[\phi^{(4)}(\langle x, \beta \rangle) \right] \Sigma \beta \beta^T \Sigma \right\|_2, \\ & \leq \|\Sigma\|_2 \sup_{\{\beta \in B_p(R)\}} \left| \mathbb{E}[\phi^{(2)}(\langle z, \beta \rangle)] - \mathbb{E}[\phi^{(2)}(\langle x, \beta \rangle)] \right| \\ & \quad + \|\Sigma\|_2^2 R^2 \sup_{\{\beta \in B_p(R)\}} \left| \mathbb{E}[\phi^{(4)}(\langle z, \beta \rangle)] - \mathbb{E}[\phi^{(4)}(\langle x, \beta \rangle)] \right|, \\ & \leq \|\Sigma\|_2 d_{\mathcal{H}_1}(x, z) + \|\Sigma\|_2^2 R^2 d_{\mathcal{H}_2}(x, z). \end{aligned}$$

Hence, we conclude that

$$\left\| \mathbb{E}[xx^T \phi^{(2)}(\langle x, \hat{\beta}^t \rangle)] - \mathfrak{E}(\hat{\beta}^t) \right\|_2 \leq d_{\mathcal{H}_3}(x, z) + \|\Sigma\|_2 d_{\mathcal{H}_1}(x, z) + \|\Sigma\|_2^2 R^2 d_{\mathcal{H}_2}(x, z).$$

□

Lemma D.3. *There exists constants c_1, c_2, c_3 depending on the eigenvalues of Σ , B , L and R such that, with probability at least $1 - c_2 e^{-c_3 p}$*

$$\left\| \frac{1}{n} \sum_{i=1}^n x_i x_i^T \int_0^1 \phi^{(2)}(\langle x_i, \beta_* + \xi(\hat{\beta}^t - \beta_*) \rangle) d\xi - \mathbb{E} \left[x x^T \int_0^1 \phi^{(2)}(\langle x, \beta_* + \xi(\hat{\beta}^t - \beta_*) \rangle) d\xi \right] \right\|_2 \leq \delta,$$

where $\delta = c_1 \sqrt{\frac{p}{n^{0.2}} \log(n)}$ for sub-gaussian covariates, and $\delta = c_1 \sqrt{\frac{p}{n} \log(n)}$ for covariates with bounded support.

Proof. We provide the proof for bounded support case. The proof for sub-gaussian case can be carried by replacing Lemma C.2 with Lemma C.5.

By the Fubini's theorem, we have

$$\begin{aligned} & \left\| \frac{1}{n} \sum_{i=1}^n x_i x_i^T \int_0^1 \phi^{(2)}(\langle x_i, \beta_* + \xi(\hat{\beta}^t - \beta_*) \rangle) d\xi - \mathbb{E} \left[x x^T \int_0^1 \phi^{(2)}(\langle x, \beta_* + \xi(\hat{\beta}^t - \beta_*) \rangle) d\xi \right] \right\|_2, \\ &= \left\| \int_0^1 \left\{ \frac{1}{n} \sum_{i=1}^n x_i x_i^T \phi^{(2)}(\langle x_i, \beta_* + \xi(\hat{\beta}^t - \beta_*) \rangle) - \mathbb{E} \left[x x^T \phi^{(2)}(\langle x, \beta_* + \xi(\hat{\beta}^t - \beta_*) \rangle) \right] \right\} d\xi \right\|_2, \\ &\leq \int_0^1 \left\| \left\{ \frac{1}{n} \sum_{i=1}^n x_i x_i^T \phi^{(2)}(\langle x_i, \beta_* + \xi(\hat{\beta}^t - \beta_*) \rangle) - \mathbb{E} \left[x x^T \phi^{(2)}(\langle x, \beta_* + \xi(\hat{\beta}^t - \beta_*) \rangle) \right] \right\} \right\|_2 d\xi, \\ &\leq \sup_{\beta \in B_p(R)} \left\| \frac{1}{n} \sum_{i=1}^n x_i x_i^T \phi^{(2)}(\langle x_i, \beta \rangle) - \mathbb{E} \left[x x^T \phi^{(2)}(\langle x, \beta \rangle) \right] \right\|_2. \end{aligned}$$

Using the definition of *operator norm*, the right hand side is equal to

$$\begin{aligned} & \sup_{\beta \in B_p(R)} \left\| \frac{1}{n} \sum_{i=1}^n x_i x_i^T \phi^{(2)}(\langle x_i, \beta \rangle) - \mathbb{E} \left[x x^T \phi^{(2)}(\langle x, \beta \rangle) \right] \right\|_2 \\ &= \sup_{\beta \in B_p(R)} \sup_{v \in S^{p-1}} \left| \frac{1}{n} \sum_{i=1}^n \phi^{(2)}(\langle x_i, \beta \rangle) \langle x_i, v \rangle^2 - \mathbb{E} \left[\phi^{(2)}(\langle x, \beta \rangle) \langle x, v \rangle^2 \right] \right|, \end{aligned}$$

where S^{p-1} denotes the p -dimensional unit sphere.

For $\Delta = 0.25$, let T_Δ be an Δ -net over S^{p-1} . Using Lemma G.4, we obtain

$$\begin{aligned}
& \mathbb{P} \left(\sup_{\beta \in B_p(R)} \sup_{v \in S^{p-1}} \left| \frac{1}{n} \sum_{i=1}^n \phi^{(2)}(\langle x_i, \beta \rangle) \langle x_i, v \rangle^2 - \mathbb{E} \left[\phi^{(2)}(\langle x, \beta \rangle) \langle x, v \rangle^2 \right] \right| > \epsilon \right), \\
& \leq \mathbb{P} \left(\sup_{\beta \in B_p(R)} \sup_{v \in T_\Delta} \left| \frac{1}{n} \sum_{i=1}^n \phi^{(2)}(\langle x_i, \beta \rangle) \langle x_i, v \rangle^2 - \mathbb{E} \left[\phi^{(2)}(\langle x, \beta \rangle) \langle x, v \rangle^2 \right] \right| > \epsilon/2 \right), \\
& \leq |T_\Delta| \mathbb{P} \left(\sup_{\beta \in B_p(R)} \left| \frac{1}{n} \sum_{i=1}^n \phi^{(2)}(\langle x_i, \beta \rangle) \langle x_i, v \rangle^2 - \mathbb{E} \left[\phi^{(2)}(\langle x, \beta \rangle) \langle x, v \rangle^2 \right] \right| > \epsilon/2 \right), \\
& = 9^p \mathbb{P} \left(\sup_{\beta \in B_p(R)} \left| \frac{1}{n} \sum_{i=1}^n \phi^{(2)}(\langle x_i, \beta \rangle) \langle x_i, v \rangle^2 - \mathbb{E} \left[\phi^{(2)}(\langle x, \beta \rangle) \langle x, v \rangle^2 \right] \right| > \epsilon/2 \right).
\end{aligned}$$

By applying Lemma C.2 to the last line above, there exists absolute constants c'_1, c'_2, c'_3 depending on L, B, R, K such that, we have

$$\mathbb{P} \left(\sup_{\beta \in B_p(R)} \left| \frac{1}{n} \sum_{i=1}^n \phi^{(2)}(\langle x_i, \beta \rangle) \langle x_i, v \rangle^2 - \mathbb{E}[\phi^{(2)}(\langle x, \beta \rangle) \langle x, v \rangle^2] \right| > c'_1 \sqrt{\frac{p}{n} \log(n)} \right) \leq c'_2 e^{-c'_3 p}.$$

c'_3 is of order $\mathcal{O}(p \log \log(n))$. Therefore, by choosing n large enough, we obtain that there exists constants c_1, c_2, c_3 such that with probability at least $1 - c_2 e^{-c_3 p}$

$$\sup_{\beta \in B} \left\| \frac{1}{n} \sum_{i=1}^n x_i x_i^T \phi^{(2)}(\langle x_i, \beta \rangle) - \mathbb{E} \left[x x^T \phi^{(2)}(\langle x, \beta \rangle) \right] \right\|_2 \leq c_1 \sqrt{\frac{p}{n} \log(n)}$$

□

Lemma D.4. *There exists a constant C depending on K and L such that,*

$$\left\| \mathbb{E}[x x^T \phi^{(2)}(\langle x, \hat{\beta}^t \rangle)] - \mathbb{E} \left[x x^T \int_0^1 \phi^{(2)}(\langle x, \beta_* + \xi(\hat{\beta}^t - \beta_*) \rangle) d\xi \right] \right\|_2 \leq \tilde{C} \|\hat{\beta}^t - \beta_*\|_2,$$

where $\tilde{C} = C$ for the bounded support case and $\tilde{C} = Cp^{1.5}$ for the sub-gaussian case.

Proof. By the Fubini's theorem, we write

$$\begin{aligned}
& \left\| \mathbb{E}[x x^T \phi^{(2)}(\langle x, \hat{\beta}^t \rangle)] - \mathbb{E} \left[x x^T \int_0^1 \phi^{(2)}(\langle x, \beta_* + \xi(\hat{\beta}^t - \beta_*) \rangle) d\xi \right] \right\|_2, \\
& = \left\| \int_0^1 \mathbb{E} \left[x x^T \left\{ \phi^{(2)}(\langle x, \hat{\beta}^t \rangle) - \phi^{(2)}(\langle x, \beta_* + \xi(\hat{\beta}^t - \beta_*) \rangle) \right\} \right] d\xi \right\|_2,
\end{aligned}$$

Moving the integration out, right hand side of above equation is smaller than

$$\begin{aligned}
& \int_0^1 \left\| \mathbb{E} \left[x x^T \left\{ \phi^{(2)}(\langle x, \hat{\beta}^t \rangle) - \phi^{(2)}(\langle x, \beta_* + \xi(\hat{\beta}^t - \beta_*) \rangle) \right\} \right] \right\|_2 d\xi, \\
& \leq \int_0^1 \left\| \mathbb{E} \left[x x^T L |\langle x, (1 - \xi)(\hat{\beta}^t - \beta_*) \rangle| \right] \right\|_2 d\xi, \\
& \leq \mathbb{E} \left[\|x\|_2^3 \|\hat{\beta}^t - \beta_*\|_2 \right] L \int_0^1 (1 - \xi) d\xi, \\
& = \frac{L \mathbb{E}[\|x\|_2^3]}{2} \|\hat{\beta}^t - \beta_*\|_2.
\end{aligned}$$

We observe that, when the covariates are supported in the ball of radius \sqrt{K} , we have $\mathbb{E}[\|x\|_2^3] \leq K^{3/2}$. When they are sub-gaussian random variables with norm K , we have $\mathbb{E}[\|x\|_2^3] \leq K^3 6^{1.5} p^{1.5}$. \square

By combining above results, for sub-gaussian covariates we obtain

$$\begin{aligned}
& \left\| [\mathbf{Q}^t]^{-1} - \int_0^1 \nabla_{\beta}^2 \ell(\beta_* + \xi(\hat{\beta}^t - \beta_*)) d\xi \right\|_2 \\
& \leq \mathfrak{D}(x, z) + c_1 \sqrt{\frac{p}{\min\{|S|p/\log(p), n/\log(n)\}}} + c_2 \|\hat{\beta}^t - \beta_*\|_2,
\end{aligned}$$

and for covariates with bounded support, we obtain

$$\begin{aligned}
& \left\| [\mathbf{Q}^t]^{-1} - \int_0^1 \nabla_{\beta}^2 \ell(\beta_* + \xi(\hat{\beta}^t - \beta_*)) d\xi \right\|_2 \\
& \leq \mathfrak{D}(x, z) + c_1 \sqrt{\frac{p}{\min\{|S|, n/\log(n)\}}} + c_2 p^{1.5} \|\hat{\beta}^t - \beta_*\|_2,
\end{aligned}$$

where

$$\mathfrak{D}(x, z) = d_{\mathcal{H}_3}(x, z) + \|\Sigma\|_2 d_{\mathcal{H}_1}(x, z) + \|\Sigma\|_2^2 R^2 d_{\mathcal{H}_2}(x, z).$$

In the following, we will derive an upper bound for $\|\mathbf{Q}^t\|_2$ where,

$$\mathbf{Q}^t = \frac{1}{\hat{\mu}_2(\hat{\beta}^t)} \left[\zeta_r(\widehat{\Sigma}_S)^{-1} - \frac{\hat{\beta}^t [\hat{\beta}^t]^T}{\hat{\mu}_2(\hat{\beta}^t)/\hat{\mu}_4(\hat{\beta}^t) + \langle \zeta_r(\widehat{\Sigma}_S) \hat{\beta}^t, \hat{\beta}^t \rangle} \right].$$

We define

$$c_L = \inf_{\beta \in B_p(L)} \mu_2(\beta).$$

Thus, for any iterate $\hat{\beta}^t$ of Newton-Stein algorithm

$$\mu_2(\hat{\beta}^t) \geq c_R.$$

By Lemmas C.1 and C.3, for some constants c_1, c_2, c_3 , with probability $1 - c_2 e^{-c_3 p}$,

$$\begin{aligned}\hat{\mu}_2(\hat{\beta}^t) &\geq \mu_2(\hat{\beta}^t) - c_1 \sqrt{\frac{p \log(n)}{n}}, \\ &\geq c_R - c_1 \sqrt{\frac{p \log(n)}{n}}.\end{aligned}$$

Also, by the assumption given in the theorem, on the set \mathcal{E} we have almost surely,

$$\inf_{t \geq 0} \left| \mu_2(\hat{\beta}^t) + \mu_4(\hat{\beta}^t) \langle \mathbf{\Sigma} \hat{\beta}^t, \hat{\beta}^t \rangle \right| > \xi,$$

for some $\xi > 0$. With probability at least $1 - c_2 e^{-c_3 p}$,

$$\begin{aligned}\left| \hat{\mu}_2(\hat{\beta}^t) + \hat{\mu}_4(\hat{\beta}^t) \langle \zeta_r(\widehat{\mathbf{\Sigma}}_S) \hat{\beta}^t, \hat{\beta}^t \rangle \right| &\geq \left| \mu_2(\hat{\beta}^t) + \mu_4(\hat{\beta}^t) \langle \mathbf{\Sigma} \hat{\beta}^t, \hat{\beta}^t \rangle \right| - \left\{ \left| \hat{\mu}_2(\hat{\beta}^t) - \mu_2(\hat{\beta}^t) \right| \right. \\ &\quad + \left| \mu_4(\hat{\beta}^t) \langle \mathbf{\Sigma} \hat{\beta}^t, \hat{\beta}^t \rangle - \hat{\mu}_4(\hat{\beta}^t) \langle \mathbf{\Sigma} \hat{\beta}^t, \hat{\beta}^t \rangle \right| \\ &\quad \left. + \left| \hat{\mu}_4(\hat{\beta}^t) \langle \mathbf{\Sigma} \hat{\beta}^t, \hat{\beta}^t \rangle - \hat{\mu}_4(\hat{\beta}^t) \langle \zeta_r(\widehat{\mathbf{\Sigma}}_S) \hat{\beta}^t, \hat{\beta}^t \rangle \right| \right\}.\end{aligned}$$

Since, we have

$$\sqrt{\frac{p}{\min\{n/\log(n), p/\log(p)|S|\}}} \leq \sqrt{\frac{p}{\min\{n/\log(n), |S|\}}},$$

by the Lemmas B.1-B.3 and Lemmas C.1-C.3, we write

$$\begin{aligned}\left| \hat{\mu}_2(\hat{\beta}^t) + \hat{\mu}_4(\hat{\beta}^t) \langle \zeta_r(\widehat{\mathbf{\Sigma}}_S) \hat{\beta}^t, \hat{\beta}^t \rangle \right| &\geq \left| \mu_2(\hat{\beta}^t) + \mu_4(\hat{\beta}^t) \langle \mathbf{\Sigma} \hat{\beta}^t, \hat{\beta}^t \rangle \right| - \left(c_1 + \|\hat{\beta}^t\|_2^2 \|\mathbf{\Sigma}\|_2 \right) \sqrt{\frac{p \log(n)}{n}} \\ &\quad - B_4 \|\hat{\beta}^t\|_2^2 \left\| \zeta_r(\widehat{\mathbf{\Sigma}}_S) - \mathbf{\Sigma} \right\|_2, \\ &\geq \left| \mu_2(\hat{\beta}^t) + \mu_4(\hat{\beta}^t) \langle \mathbf{\Sigma} \hat{\beta}^t, \hat{\beta}^t \rangle \right| - C \sqrt{\frac{p}{\min\{n/\log(n), |S|\}}}, \\ &\geq \xi - C \sqrt{\frac{p}{\min\{n/\log(n), |S|\}}},\end{aligned}$$

where $C = \max\{cB_4R^2, c_1 + R^2\|\mathbf{\Sigma}\|_2\}$.

Therefore, for some constants c_1, c_2 , with high probability, we have

$$\begin{aligned}\|\mathbf{Q}^t\|_2 &\leq \frac{1}{\hat{\mu}_2(\hat{\beta}^t)} \left[\left\| \zeta_r(\widehat{\mathbf{\Sigma}}_S)^{-1} \right\|_2 + \frac{|\hat{\mu}_4(\hat{\beta}^t)| \|\hat{\beta}^t\|_2^2}{\left| \hat{\mu}_2(\hat{\beta}^t) + \hat{\mu}_4(\hat{\beta}^t) \langle \zeta_r(\widehat{\mathbf{\Sigma}}_S) \hat{\beta}^t, \hat{\beta}^t \rangle \right|} \right], \\ &\leq \frac{1}{c_R - c_1 \sqrt{\frac{p \log(n)}{n}}} \left[\frac{1}{\hat{\sigma}^2} + \frac{B_4 R^2}{\xi - C \sqrt{\frac{p}{\min\{n/\log(n), |S|\}}}} \right], \\ &\leq \frac{1}{c_R - c_1 \sqrt{\frac{p \log(n)}{n}}} \left[\frac{1}{\sigma^2 - c_2 \sqrt{\frac{\log(p)}{|S|}}} + \frac{B_4 R^2}{\xi - C \sqrt{\frac{p}{\min\{n/\log(n), |S|\}}}} \right],\end{aligned}$$

For n and $|S|$ sufficiently large so that we have the following inequalities,

$$\begin{aligned} c_2 \sqrt{\frac{\log(p)}{|S|}} &\leq \frac{\sigma^2}{2}, \\ c_1 \sqrt{\frac{p \log(n)}{n}} &\leq \frac{c_R}{2}, \\ C \sqrt{\frac{p}{\min\{n/\log(n), |S|\}}} &\leq \frac{\xi}{2}, \end{aligned}$$

we obtain

$$\|\mathbf{Q}^t\|_2 \leq \frac{2}{c_R} \left[\frac{2}{\sigma^2} + \frac{2B_4 R^2}{\xi} \right] := \kappa.$$

Finally, we take into account the conditioning on the event \mathcal{E} and conclude the proof.

Proof of Corollaries 4.2 and 4.5. In the following, we provide the proof for Corollary 4.2. The proof for Corollary 4.5 follows from the exact same steps.

The statement of the Theorem 4.1 holds on the probability space with a probability lower bounded by $\mathbb{P}(\mathcal{E}) - c/p^2$ for some constant c . Let \mathcal{Q} denote this set, on which the statement of the lemma holds. Note that $\mathcal{Q} \subset \mathcal{E}$. We have

$$\mathbb{P}(\mathcal{Q}) \geq \mathbb{P}(\mathcal{E}) - c'/p^2.$$

This suggests that the difference between \mathcal{Q} and \mathcal{E} is small. By taking expectations on both sides over the set \mathcal{Q} , we obtain,

$$\begin{aligned} \mathbb{E} \left[\|\hat{\beta}^{t+1} - \beta_*\|_2; \mathcal{Q} \right] &\leq \kappa \left\{ \mathfrak{D}(x, z) + c_1 \sqrt{\frac{p}{\min\{p/\log(p)|S|, n/\log(n)\}}} \right\} \mathbb{E} \left[\|\hat{\beta}^t - \beta_*\|_2 \right] \\ &\quad + \kappa c_2 \mathbb{E} \left[\|\hat{\beta}^t - \beta_*\|_2^2 \right] \end{aligned}$$

where we used

$$\mathbb{E} \left[\|\hat{\beta}^t - \beta_*\|_2^l; \mathcal{Q} \right] \leq \mathbb{E} \left[\|\hat{\beta}^t - \beta_*\|_2^l \right], \quad l = 1, 2.$$

Similarly for the iterate $\hat{\beta}^{t+1}$, we write

$$\begin{aligned} \mathbb{E} \left[\|\hat{\beta}^{t+1} - \beta_*\|_2 \right] &= \mathbb{E} \left[\|\hat{\beta}^{t+1} - \beta_*\|_2; \mathcal{Q} \right] + \mathbb{E} \left[\|\hat{\beta}^{t+1} - \beta_*\|_2; \mathcal{Q}^C \right], \\ &\leq \mathbb{E} \left[\|\hat{\beta}^{t+1} - \beta_*\|_2; \mathcal{Q} \right] + 2R\mathbb{P}(\mathcal{Q}^C), \\ &\leq \mathbb{E} \left[\|\hat{\beta}^{t+1} - \beta_*\|_2; \mathcal{Q} \right] + 2R \left(\mathbb{P}(\mathcal{E}^C) + \frac{c}{p^2} \right), \\ &\leq \mathbb{E} \left[\|\hat{\beta}^{t+1} - \beta_*\|_2; \mathcal{Q} \right] + \frac{\epsilon}{10}, \\ &\leq \mathbb{E} \left[\|\hat{\beta}^{t+1} - \beta_*\|_2; \mathcal{Q} \right] + \frac{\mathbb{E} \left[\|\hat{\beta}^t - \beta_*\|_2 \right]}{10}. \end{aligned}$$

Combining these two inequalities, we obtain

$$\begin{aligned} \mathbb{E} \left[\|\hat{\beta}^{t+1} - \beta_*\|_2 \right] &\leq \left\{ 0.1 + \kappa \mathfrak{D}(x, z) + c_1 \kappa \sqrt{\frac{p}{\min \{p/\log(p)|S|, n/\log(n)\}}} \right\} \mathbb{E} \left[\|\hat{\beta}^t - \beta_*\|_2 \right] \\ &\quad + c_2 \kappa \mathbb{E} \left[\|\hat{\beta}^t - \beta_*\|_2^2 \right]. \end{aligned}$$

Hence the proof follows. \square

Proof of Theorem 4.3. For a sequence satisfying the following inequality,

$$\|\hat{\beta}^{t+1} - \beta_*\|_2 \leq \left(\tau_1 + \tau_2 \|\hat{\beta}^t - \beta_*\|_2 \right) \|\hat{\beta}^t - \beta_*\|_2,$$

we observe that

$$\tau_1 + \tau_2 \|\hat{\beta}^0 - \beta_*\|_2 < 1 \tag{D.2}$$

is a sufficient condition for convergence to 0. Let $\xi \in (\epsilon, 1)$ and t_ξ be the last iteration that $\|\hat{\beta}^t - \beta_*\|_2 > \delta$. Then, for $t > t_\xi$

$$\begin{aligned} \|\hat{\beta}^{t+1} - \beta_*\|_2 &\leq \left(\tau_1 + \tau_2 \|\hat{\beta}^t - \beta_*\|_2 \right) \|\hat{\beta}^t - \beta_*\|_2, \\ &\leq (\tau_1 + \tau_2 \xi) \|\hat{\beta}^t - \beta_*\|_2. \end{aligned}$$

This convergence behavior describes a linear rate and requires at most

$$\frac{\log(\epsilon/\xi)}{\log(\tau_1 + \tau_2 \xi)}$$

iterations to reach a tolerance of ϵ . For $t \leq t_\xi$, we have

$$\begin{aligned} \|\hat{\beta}^{t+1} - \beta_*\|_2 &\leq \left(\tau_1 + \tau_2 \|\hat{\beta}^t - \beta_*\|_2 \right) \|\hat{\beta}^t - \beta_*\|_2, \\ &\leq (\tau_1/\xi + \tau_2) \|\hat{\beta}^t - \beta_*\|_2^2. \end{aligned}$$

This describes a quadratic rate and the number of iterations to reach a tolerance of ξ can be upper bounded by

$$\log_2 \left(\frac{\log(\delta(\tau_1/\xi + \tau_2))}{\log((\tau_1/\xi + \tau_2) \|\hat{\beta}^0 - \beta_*\|_2)} \right).$$

Therefore, the overall number of iterations to reach a tolerance of ϵ is upper bounded by

$$\mathcal{J}(\xi) = \log_2 \left(\frac{\log(\delta(\tau_1/\xi + \tau_2))}{\log((\tau_1/\xi + \tau_2) \|\hat{\beta}^0 - \beta_*\|_2)} \right) + \frac{\log(\epsilon/\xi)}{\log(\tau_1 + \tau_2 \xi)}$$

which is a function of ξ . Therefore, we take the minimum over the feasible set. \square

E Step size selection

We carry our analysis from Eq. D.1. The optimal step size would be

$$\gamma_* = \operatorname{argmin}_{\gamma} \left\| I - \gamma \mathbf{Q}^t \int_0^1 \nabla_{\beta}^2 \ell(\beta_* + \xi(\hat{\beta}^t - \beta_*)) d\xi \right\|_2.$$

Using the mean value theorem,

$$\nabla_{\beta}^2 \ell(\tilde{\beta}) = \int_0^1 \nabla_{\beta}^2 \ell(\beta_* + \xi(\hat{\beta}^t - \beta_*)) d\xi,$$

where $\tilde{\beta} \in [\beta_*, \hat{\beta}^t]$, and write the governing term as

$$\left\| I - \gamma \mathbf{Q}^t \nabla_{\beta}^2 \ell(\tilde{\beta}) \right\|_2.$$

The above function is piecewise linear in γ and it can be minimized by setting

$$\gamma_* = \frac{2}{\lambda_1 \left(\mathbf{Q}^t \nabla_{\beta}^2 \ell(\tilde{\beta}) \right) + \lambda_p \left(\mathbf{Q}^t \nabla_{\beta}^2 \ell(\tilde{\beta}) \right)}.$$

Since we don't have access to the optimal value β_* nor $\tilde{\beta}$, we will assume that $\tilde{\beta}$ and the current iterate $\hat{\beta}^t$ are close.

In the regime $n \gg p$, and by our construction of the scaling matrix \mathbf{Q}^t , we have

$$\mathbf{Q}^t \approx \left[\mathbb{E}[xx^T \phi^{(2)}(\langle x, \hat{\beta}^t \rangle)] \right]^{-1} \quad \text{and} \quad \nabla_{\beta}^2 \ell(\hat{\beta}^t) \approx \left[\mathbb{E}[xx^T \phi^{(2)}(\langle x, \hat{\beta}^t \rangle)] \right].$$

The crucial observation is that ζ_r function sets the smallest eigenvalue to $\hat{\sigma}^2$ which overestimates σ^2 in general. Even though the largest eigenvalue of $\mathbf{Q}^t \nabla_{\beta}^2 \ell(\tilde{\beta})$ will be close to 1, the smallest value will be $\sigma^2/\hat{\sigma}^2$. This will make the optimal step size larger than 1. Hence, we suggest

$$\gamma = \frac{2}{1 + \sigma^2/\hat{\sigma}^2}.$$

We also have, by the Weyl's inequality,

$$|\hat{\sigma}^2 - \sigma^2| \leq \left\| \hat{\Sigma} - \Sigma \right\|_2 \leq C \sqrt{\frac{p}{|S|}},$$

with high probability. Whenever r is less than $p/2$, we suggest

$$\gamma = \frac{2}{1 + \frac{\hat{\sigma}^2 - \mathcal{O}(\sqrt{p/|S|})}{\hat{\sigma}^2}}.$$

F Details of experiments

Table 2 provides the details of the experiments given in Section 5.

S3					S20				
	LR		LS			LR		LS	
Method	Elapsed(sec)	Iter	Elapsed(sec)	Iter	Method	Elapsed(sec)	Iter	Elapsed(sec)	Iter
NewSt	10.637	2	8.763	4	NewSt	23.158	4	16.475	10
BFGS	22.885	8	13.149	6	BFGS	40.258	17	54.294	37
LBFGS	46.763	19	19.952	11	LBFGS	51.888	26	33.107	20
Newton	55.328	2	38.150	1	Newton	47.955	2	39.328	1
GD	865.119	493	155.155	100	GD	1204.015	245	145.987	100
AGD	169.473	82	65.396	42	AGD	182.031	83	56.257	38

CT Slices					Coertype				
	LR		LS			LR		LS	
Method	Elapsed(sec)	Iter	Elapsed(sec)	Iter	Method	Elapsed(sec)	Iter	Elapsed(sec)	Iter
NewSt	4.191	32	1.799	11	NewSt	16.113	31	2.080	5
BFGS	4.638	35	4.525	37	BFGS	21.916	48	2.238	3
LBFGS	26.838	217	22.679	180	LBFGS	30.765	69	2.321	3
Newton	5.730	3	1.937	1	Newton	122.158	40	2.164	1
GD	96.142	1156	61.526	721	GD	194.473	446	22.738	60
AGD	96.142	880	45.864	518	AGD	80.874	186	32.563	77

Table 2: Details of the experiments presented in Figures 2 and 3.

G Useful lemmas

Lemma G.1 (Gautschi’s Inequality). *Let Γ denote the Gamma function. Then, for $r \in (0, 1)$, we have*

$$z^{1-r} < \frac{\Gamma(z+1)}{\Gamma(z+r)} < (1+z)^{1-r}.$$

Lemma G.2. *Let Z be a random variable with a density function f and cumulative distribution function F . If $F^C = 1 - F$, then,*

$$|\mathbb{E}[Z \mathbb{1}_{\{|Z|>t\}}]| \leq t\mathbb{P}(|Z| > t) + \int_t^\infty \mathbb{P}(|Z| > z)dz.$$

Proof. We write,

$$\mathbb{E}[Z \mathbb{1}_{\{|Z|>t\}}] = \int_t^\infty zf(z)dz + \int_{-\infty}^{-t} zf(z)dz.$$

Using integration by parts, we obtain

$$\begin{aligned} \int zf(z)dz &= -zF^C(z) + \int F^C(z)dz, \\ &= zF(z) - \int F(z)dz. \end{aligned}$$

Since $\lim_{z \rightarrow \infty} zF^C(z) = \lim_{z \rightarrow -\infty} zF(z) = 0$, we have

$$\begin{aligned} \int_t^\infty z f(z) dz &= tF^C(t) + \int_t^\infty F^C(z) dz, \\ \int_{-\infty}^{-t} z f(z) dz &= -tF(-t) - \int_{-\infty}^{-t} F(z) dz, \\ &= -tF(-t) - \int_t^\infty F(-z) dz. \end{aligned}$$

Hence, we obtain the following bound,

$$\begin{aligned} |\mathbb{E}[Z \mathbb{1}_{\{|Z| > t\}}]| &= \left| tF^C(t) + \int_t^\infty F^C(z) dz - tF(-t) - \int_t^\infty F(-z) dz \right|, \\ &\leq t(F^C(t) + F(-t)) + \left(\int_t^\infty F^C(z) + F(-z) dz \right), \\ &\leq t\mathbb{P}(|Z| > t) + \int_t^\infty \mathbb{P}(|Z| > z) dz. \end{aligned}$$

□

Lemma G.3. *For positive constants c_1, c_2 , we have*

$$\int_{c_1}^\infty e^{-c_2 t^{2/3}} dt \leq \left\{ \frac{3c_1^{1/3}}{2c_2} + \frac{3}{4c_2^2 c_1^{1/3}} \right\} e^{-c_2 c_1^{2/3}}$$

Proof. By the change of variables $t^{2/3} = x^2$, we get

$$\int_{c_1}^\infty e^{-c_2 t^{2/3}} dt = 3 \int_{c_1^{1/3}}^\infty x^2 e^{-c_2 x^2} dx.$$

Next, we notice that

$$de^{-c_2 x^2} = -2c_2 x e^{-c_2 x^2} dx.$$

Hence, using the integration by parts, we have

$$\int_{c_1}^\infty e^{-c_2 t^{2/3}} dt = \frac{3}{2c_2} \left\{ c_1^{1/3} e^{-c_2 c_1^{2/3}} + \int_{c_1^{1/3}}^\infty e^{-c_2 x^2} dx \right\}.$$

We will find an upper bound on the second term. Using the change of variables, $x = y + c_1^{1/3}$, we obtain

$$\begin{aligned} \int_{c_1^{1/3}}^\infty e^{-c_2 x^2} dx &= \int_0^\infty e^{-c_2 (y + c_1^{1/3})^2} dy, \\ &\leq e^{-c_2 c_1^{2/3}} \int_0^\infty e^{-2c_2 y c_1^{1/3}} dy, \\ &= \frac{e^{-c_2 c_1^{2/3}}}{2c_2 c_1^{1/3}}. \end{aligned}$$

Combining the above results, we complete the proof. □

Lemma G.4 ([Ver10]). *Let X be a symmetric $p \times p$ matrix, and let T_ϵ be an ϵ -net over S^{p-1} . Then,*

$$\|X\|_2 \leq \frac{1}{1-2\epsilon} \sup_{v \in T_\epsilon} |\langle Xv, v \rangle|.$$

Lemma G.5. *Let $B_p(R) \subset \mathbb{R}^p$ be the ball of radius R centered at the origin and T_ϵ be an ϵ -net over $B_p(R)$. Then,*

$$|T_\epsilon| \leq \left(\frac{R\sqrt{p}}{\epsilon} \right)^p.$$

Proof of Lemma G.5. A similar proof appears in [VdV00]. The set $B_p(R)$ can be contained in a p -dimensional cube of size $2R$. Consider a grid over this cube with mesh width $2\epsilon/\sqrt{p}$. Then $B_p(R)$ can be covered with at most $(2R/(2\epsilon/\sqrt{p}))^p$ many cubes of edge length $2\epsilon/\sqrt{p}$. If one takes the projection of the centers of such cubes onto $B_p(R)$ and considers the circumscribed balls of radius ϵ , we may conclude that $B_p(R)$ can be covered with at most

$$\left(\frac{2R}{2\epsilon/\sqrt{p}} \right)^p$$

many balls of radius ϵ . □

Lemma G.6. *For $a, b > 0$, and ϵ satisfying*

$$\epsilon = \left\{ \frac{a}{2} \log \left(\frac{2b^2}{a} \right) \right\}^{1/2} \quad \text{and} \quad \frac{2}{a} b^2 > e,$$

we have $\epsilon^2 \geq a \log(b/\epsilon)$.

Proof. Since $a, b > 0$ and $x \rightarrow e^x$ is a monotone increasing function, the above inequality condition is equivalent to

$$\frac{2\epsilon^2}{a} e^{\frac{2\epsilon^2}{a}} \geq \frac{2b^2}{a}.$$

Now, we define the function $f(w) = we^w$ for $w > 0$. f is continuous and invertible on $[0, \infty)$. Note that f^{-1} is also a continuous and increasing function for $w > 0$. Therefore, we have

$$\epsilon^2 \geq \frac{a}{2} f^{-1} \left(\frac{2b^2}{a} \right)$$

Observe that the smallest possible value for ϵ would be simply the square root of $af^{-1}(2b^2/a)/2$. For simplicity, we will obtain a more interpretable expression for ϵ . By the definition of f^{-1} , we have

$$\log(f^{-1}(y)) + f^{-1}(y) = \log(y).$$

Since the condition on a and b enforces $f^{-1}(y)$ to be larger than 1, we obtain the simple inequality that

$$f^{-1}(y) \leq \log(y).$$

Using the above inequality, if ϵ satisfies

$$\epsilon^2 = \frac{a}{2} \log \left(\frac{2b^2}{a} \right) \geq \frac{a}{2} g \left(\frac{2b^2}{a} \right),$$

we obtain the desired inequality. □